

VERITAS NetBackup™ Vault 5.1

System Administrator's Guide

for UNIX and Windows

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2001–2004 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS logo, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS, NetBackup, the VERITAS logo, Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

Portions of this software are derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm. Copyright 1991-92, RSA Data Security, Inc. Created 1991. All rights reserved.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650–527–8000 Fax 650–527–2908
www.veritas.com

Third-Party Copyrights

ACE 5.2A: ACE(TM) is copyrighted by Douglas C.Schmidt and his research group at Washington University and University of California, Irvine, Copyright (c) 1993-2002, all rights reserved.

IBM XML for C++ (XML4C) 3.5.1: Copyright (c) 1999,2000,2001 Compaq Computer Corporation; Copyright (c) 1999,2000,2001 Hewlett-Packard Company; Copyright (c) 1999,2000,2001 IBM Corporation; Copyright (c) 1999,2000,2001 Hummingbird Communications Ltd.; Copyright (c) 1999,2000,2001 Silicon Graphics, Inc.; Copyright (c) 1999,2000,2001 Sun Microsystems, Inc.; Copyright (c) 1999,2000,2001 The Open Group; All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

JacORB 1.4.1: The licensed software is covered by the GNU Library General Public License, Version 2, June 1991.

Open SSL 0.9.6: This product includes software developed by the OpenSSL Project * for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)

TAO (ACE ORB) 1.2a: TAO(TM) is copyrighted by Douglas C. Schmidt and his research group at Washington University and University of California, Irvine, Copyright (c) 1993-2002, all rights reserved.



Contents

| | |
|--|-----------|
| Preface | xv |
| What Is In This Guide? | xv |
| Getting Help | xvi |
| NetBackup Vault Manuals | xviii |
| Glossary | xix |
| Accessibility Features | xx |
| Conventions | xx |
| Chapter 1. Introduction to Vault | 1 |
| Vault Uses NetBackup Functions | 1 |
| How to Access NetBackup Vault | 2 |
| Vault Original or Duplicate Images? | 2 |
| The Vault Process | 2 |
| Choose Backup Images | 3 |
| Duplicate Backup Images | 3 |
| Backup the NetBackup Catalog | 4 |
| Eject Media | 4 |
| Generate Reports | 4 |
| How Vault Uses Volume Groups and Pools | 4 |
| NetBackup and Vault Configuration | 5 |
| Chapter 2. Installing NetBackup Vault | 7 |
| Supported Systems | 7 |
| Supported Robots | 7 |



| | |
|---|-----------|
| UNIX Systems | 7 |
| Installation Prerequisites for a UNIX System | 8 |
| Installing NetBackup Vault on a UNIX System | 8 |
| Upgrading NetBackup Vault on a UNIX System | 10 |
| Uninstalling NetBackup Vault from a UNIX System | 13 |
| Microsoft Windows Systems | 15 |
| Licensing Prerequisites for a Windows System | 16 |
| Licensing NetBackup Vault on a Windows System | 16 |
| Upgrading NetBackup Vault on a Windows System | 17 |
| Delicensing NetBackup Vault from a Windows System | 17 |
| Upgrading from bpvault 3.4 | 18 |
| Chapter 3. Best Practices | 19 |
| Vaulting Paradigm | 20 |
| Preferred Vaulting Strategies | 20 |
| Vault Original Backups | 21 |
| Use Disk Staging | 21 |
| Ensure All Data is Vaulted | 22 |
| Overlap the Time Window in the Profile | 22 |
| Consequences of Not Overlapping the Time Window: Missing Data | 23 |
| Resolve Multiple Names for a Single Server | 23 |
| Specify Robotic Volume Group When Configuring a Vault | 23 |
| Multiple Volume Groups (Multiple Robots) | 23 |
| Do Not Vault More Than You Need To | 24 |
| Send Only the Intended Backups Off-site | 24 |
| Avoid Vaulting Partial Images | 24 |
| Vaulting Original Backups in a 24x7 Environment | 25 |
| Preparing for Efficient Recovery | 26 |
| Vault NetBackup Catalogs | 26 |
| Use Precise Naming Conventions for Volume Pools and Groups | 27 |



| | |
|---|----|
| Match Volume Pools to Data Usage | 27 |
| Designate a Primary Copy and Keep It On Site | 27 |
| Suspend Vaulted Media | 28 |
| Revault Unexpired Media | 29 |
| Defer Ejection | 29 |
| Avoid Resource Contention During Duplication | 29 |
| When Two Processes Try to Use the Same Drive | 30 |
| Alternative A: Dedicated Robot for Vault Processing | 30 |
| Alternative B: Each Robot as a Vault Robot | 31 |
| Alternative C: One Robot as Both a Backup and Vault Robot | 32 |
| When the Read Drive Is Not in the Vault Robot | 33 |
| Sharing Resources with Backup Jobs | 33 |
| Load Balancing | 34 |
| Profiles for Both Originals and Duplicates | 34 |
| If Your Vault Vendor Does Not Pick Up Media Every Day | 34 |
| Specifying Different Volume Pools for Source and Destination | 35 |
| Avoid Sending Duplicates Over The Network | 35 |
| Create Originals Concurrently | 35 |
| Use Alternate Read Server | 35 |
| Use Advanced Duplication Configuration | 36 |
| Increase Duplication Throughput | 37 |
| Configuring for Multiple-Drives: Basics | 38 |
| Multiple-Drive Scenario: Does Not Send Data Over Network | 38 |
| Maximize Drive Utilization During Duplication | 39 |
| Use Scratch Volume Pools | 40 |
| Ensure Report Integrity | 40 |
| Organizing Reports by Robot | 41 |
| Organizing Reports by Vault | 41 |
| Organizing Reports by Profile | 41 |
| Consequences of Sharing an Off-site Volume Group Across Multiple Robots ... | 41 |



| | |
|---|-----------|
| Generate the Lost Media Report Regularly | 41 |
| Chapter 4. Preparing NetBackup for Vault | 43 |
| Volume Pools | 43 |
| Volume Groups | 45 |
| Vault Policies | 46 |
| Policy Configuration Information | 46 |
| Creating a Vault Policy | 47 |
| Chapter 5. Configuring Vault | 49 |
| Information Required to Configure Vault | 49 |
| Master Server, Media Servers, and Storage Units | 50 |
| Robot Information | 51 |
| Methods of Configuration | 51 |
| Configuring Vault Properties | 52 |
| E-mail Tab | 52 |
| Alternate Media Server Names Tab | 53 |
| Alternative Media Server Names Background | 54 |
| Alternate Media Server Names Considerations | 54 |
| How to Add Alternate Media Server Names | 55 |
| Configuring Robots for Vault | 56 |
| Vault Robot Dialog | 56 |
| Creating a Vault | 57 |
| Vault Dialog | 57 |
| Requirements for Creating a Vault | 58 |
| How to Create a Vault | 59 |
| Vault Dialog Configuration Options | 59 |
| Creating a Profile | 61 |
| Profile Dialog | 61 |
| How to Create a Profile | 62 |
| Configuring a Profile | 62 |



| | |
|--|-----------|
| Configuring Choose Backups | 63 |
| Choose Backups Tab | 63 |
| Choose Backups Configuration Options | 64 |
| Configuring Duplication | 66 |
| Duplication Tab | 66 |
| The Primary Backup Image | 67 |
| Basic Duplication | 67 |
| Advanced Duplication | 68 |
| Duplication Tab Configuration Options | 70 |
| Multiple Copies Dialog | 74 |
| Duplication Rule Dialog | 77 |
| Treatment of Images Without Corresponding Duplication Rule | 80 |
| Configuring Catalog Backup | 82 |
| Catalog Backup Tab | 82 |
| Default Catalog Locations | 83 |
| Catalog Backup Configuration Options | 84 |
| Configuring Eject | 85 |
| Eject Tab | 85 |
| Eject Configuration Options | 87 |
| Media Ejection Overview | 88 |
| ACS MAP Overview | 89 |
| Eject Mode (Immediate or Deferred) | 89 |
| Media Ejection Timeout Impact | 91 |
| Configuring Reports | 91 |
| Reports Tab | 92 |
| Reports Configuration Options | 93 |
| Report Mode (Immediate or Deferred) | 95 |
| Reports that Depend on Eject | 96 |
| Chapter 6. Vaulting and Managing Media | 97 |



| | |
|--|-----|
| Running a Vault Session | 98 |
| Running Multiple Sessions Simultaneously | 98 |
| Running a Session Automatically | 99 |
| Running a Session Manually | 99 |
| Running a Session from the Administration Console | 99 |
| Running a Session from a Command Line | 100 |
| Previewing a Vault Session | 100 |
| Stopping a Vault Session | 101 |
| Resuming a Vault Session | 101 |
| Monitoring a Vault Session | 102 |
| Extended Error Codes | 103 |
| The List of Images to be Vaulted | 104 |
| Duplication Exclusions | 104 |
| Ejection Exclusions | 105 |
| Vault Resiliency | 105 |
| Ejecting Media | 106 |
| Previewing Media To Be Ejected | 106 |
| Ejecting Media by Using the NetBackup Administration Console | 107 |
| Ejecting Media by Using the Vault Operator Menu | 108 |
| Ejecting Media by Using the vlteject Command | 109 |
| Ejecting Media by Using a Vault Policy | 110 |
| Consolidating Ejects | 111 |
| Injecting Media | 112 |
| Injecting Media by Using the Administration Console | 112 |
| Injecting Media by Using the Vault Operator Menu | 113 |
| Injecting Media by Using the vltinject Command | 114 |
| Vaulting and Managing Media in Containers | 115 |
| Vaulting Media in Containers | 116 |
| Vaulting Container Media by Using the Vault Operator Menu | 116 |
| Vaulting Container Media by Using the vltcontainers Command | 116 |



| | |
|---|------------|
| Managing Containers and Media | 118 |
| Using the Vault Operator Menu to Manage Container Media | 118 |
| Using the vltcontainers Command to Manage Container Media | 119 |
| Reporting on Containers and Media | 120 |
| Assigning Multiple Retentions with One Profile | 121 |
| Vaulting Additional Volumes | 124 |
| Duplicating a Volume Manually | 124 |
| Duplicating a Volume by Using Vault | 125 |
| Revaulting Unexpired Media | 126 |
| Tracking Volumes Not Ejected by Vault | 128 |
| Vaulting Media Not Created by NetBackup | 129 |
| Notifying a Tape Operator When Eject Begins | 130 |
| Using Notify Scripts | 131 |
| Notify Script for a Specific Robot | 132 |
| Notify Script for a Specific Vault | 133 |
| Notify Script for a Specific Profile | 133 |
| Order of Execution | 133 |
| Clearing the Media Description Field | 134 |
| Ensuring Available Media for Catalog Backups | 134 |
| Deassigning Vaulted NetBackup Catalog Media | 135 |
| Restoring Data from Vaulted Media | 136 |
| Replacing Damaged Media | 137 |
| Chapter 7. Creating Originals or Copies Concurrently | 143 |
| Understanding Concurrent Copies | 143 |
| Continue or Fail for Concurrent Copies | 144 |
| Continue Copies | 144 |
| Fail All Copies | 145 |
| Creating Original Images Concurrently | 145 |
| Creating Duplicate Images Concurrently | 147 |



| | |
|---|------------|
| When Duplication is Possible | 148 |
| Concurrent Copies through the Catalog Node | 149 |
| Concurrent Copies during Basic Duplication | 151 |
| Concurrent Copies during Advanced Duplication | 155 |
| Chapter 8. Reporting | 159 |
| Generating Reports | 159 |
| Generating Reports by Using the Vault Operator Menu | 160 |
| Generating Reports by Using the vlteject Command | 160 |
| Generating Reports by Using a Vault Policy | 161 |
| Consolidating Reports | 162 |
| Viewing Reports | 163 |
| Vault Report Types | 164 |
| Reports for Media Going Off-Site | 164 |
| Picking List for Robot | 164 |
| Distribution List for Vault | 165 |
| Detailed Distribution List for Vault | 166 |
| Summary Distribution List for Vault | 167 |
| Reports for Media Coming On-Site | 168 |
| Picking List for Vault | 168 |
| Distribution List for Robot | 169 |
| Inventory Reports | 170 |
| Vault Inventory | 170 |
| Off-site Inventory | 171 |
| All Media Inventory | 172 |
| Graphical Representation of Inventory Reports Scope | 173 |
| Container Inventory Report | 173 |
| Recovery Report for Vault | 174 |
| Lost Media Report | 175 |
| Non-vaulted Images Exception Report | 175 |



| | |
|---|------------|
| Iron Mountain FTP File | 176 |
| Chapter 9. Administering Vault | 179 |
| Setting Up E-Mail | 179 |
| Administering Access to Vault | 180 |
| Vault Operator User Group | 180 |
| Printing Vault and Profile Information | 182 |
| Copying a Profile | 182 |
| Moving a Vault to a Different Robot | 183 |
| Changing Volume Pools and Groups | 184 |
| Vault Session Log Files | 184 |
| Session Logs | 185 |
| Setting the Duration of Vault Session Files | 186 |
| Output from the Vault Driver | 186 |
| General Operational Issues | 187 |
| Vaulting Storage Migrator Files | 187 |
| Disk Only Source of Backups | 187 |
| The Scope of the Source Volume Group | 187 |
| Chapter 10. Using the Menu User Interfaces | 189 |
| Using the Vault Administration Interface | 189 |
| Using the Vault Operator Menu Interface | 191 |
| Changes in vmadm for Vault | 192 |
| Additions to Volume Configuration | 192 |
| Changes to the Special Actions Menu | 192 |
| Change Vault Name for Volumes | 192 |
| Change Date Volumes are Sent to Vault | 193 |
| Change Date Volumes Return from Vault | 193 |
| Change Vault Slot for Volumes | 194 |
| Change Vault Session ID for Volumes | 194 |
| Changes to Display Options | 194 |



| | |
|--|------------|
| Changes in bpdjobs for Vault | 195 |
| Chapter 11. Troubleshooting | 197 |
| Printing Problems | 197 |
| Errors Returned by the Vault Session | 198 |
| No Media Are Ejected | 198 |
| Media is Missing in Robot | 198 |
| Bad or Missing Duplicate Tape | 199 |
| Tape Drive or Robot Offline | 200 |
| No Duplicate Progress Message | 200 |
| Ejecting Tapes While in Use | 201 |
| Tapes Ejected to the MAP are Returned to Robot | 201 |
| Unexpired Tapes Were Injected into the Robot | 201 |
| Vault Session Locking | 202 |
| Debug Logs | 202 |
| NetBackup Debug Logs | 203 |
| Setting the Duration and Level of Debug Logs | 203 |
| Logs To Accompany Problem Reports | 204 |
| Appendix A. Recovering from Disasters | 205 |
| Introduction | 205 |
| Definition of Disaster | 206 |
| Definition of Disaster Recovery | 206 |
| Definition of Disaster Recovery Plan | 207 |
| Recovery Priorities | 207 |
| Developing a Disaster Recovery Plan | 207 |
| Testing a Disaster Recovery Plan | 209 |
| Disaster Recovery in the NetBackup Vault Context | 209 |
| Preparing for Recovery | 209 |
| Recovering NetBackup | 211 |
| Recovering Data | 212 |



| | |
|---|------------|
| Recovering to a Specific Point in Time | 215 |
| Appendix B. Vault's File and Directory Structure | 217 |
| UNIX Files and Directories | 217 |
| Windows Files and Directories | 224 |
| Appendix C. Vault Functional Design | 231 |
| Functional Design Overview | 231 |
| Other Related Services | 232 |
| Other Related Vault Documents | 232 |
| Architectural Services | 232 |
| Services Interactions Diagram | 233 |
| Client/Server Architectural Services | 234 |
| Technical Components | 236 |
| Components for Vault | 236 |
| Technical Design Issues | 236 |
| Vault Technical Components | 237 |
| Technical Example: Standard Duplication Diagram | 242 |
| Technical Example: Standard Duplication Table | 242 |
| Operational Procedures | 246 |
| NetBackup Vault Image Duplication Process | 247 |
| Index | 251 |





Preface

VERITAS NetBackup™ Vault simplifies the processes of image duplication, off-site storage, and off-site retrieval for both storage administrators and system operators. This *System Administrator's Guide* explains procedures performed by the system administrator and provides the administrative information needed to run NetBackup Vault on both UNIX and Windows platforms. A separate *Operator's Guide* provides instructions for system operators.

What Is In This Guide?

This guide contains information that applies to both NetBackup Server and NetBackup Enterprise Server. Information that applies to NetBackup Server or NetBackup Enterprise Server specifically is noted where appropriate.

This guide is organized as follows:

- ◆ [Chapter 1, "Introduction to Vault,"](#) introduces Vault and describes how Vault works.
- ◆ [Chapter 2, "Installing NetBackup Vault,"](#) presents the steps required to install and configure Vault.
- ◆ [Chapter 3, "Best Practices,"](#) provides advice on how to configure Vault efficiently and avoid configuration problems.
- ◆ [Chapter 4, "Preparing NetBackup for Vault,"](#) provides instructions for setting up NetBackup for use with Vault and for configuring NetBackup policies for Vault.
- ◆ [Chapter 5, "Configuring Vault,"](#) describes procedures for configuring Vault robots, vaults, and profiles.
- ◆ [Chapter 6, "Vaulting and Managing Media,"](#) describes procedures for running Vault sessions and managing vaulted media.
- ◆ [Chapter 7, "Creating Originals or Copies Concurrently,"](#) provides information about how NetBackup and Vault can create multiple backup images concurrently.
- ◆ [Chapter 8, "Reporting,"](#) details the reports available through Vault, how they are generated, and how to receive notification of Vault activity.



- ◆ [Chapter 9, “Administering Vault,”](#) provides information about administrative tasks in Vault.
- ◆ [Chapter 10, “Using the Menu User Interfaces,”](#) explains the Vault functionality available through the Vault Administration and Vault Operator Menu interfaces.
- ◆ [Chapter 11, “Troubleshooting,”](#) discusses potential problems that may occur when using Vault and how to resolve or work around them.
- ◆ [Appendix A, “Recovering from Disasters,”](#) describes how to prepare for a disaster and how to recover media after a disaster or some other event damages media.
- ◆ [Appendix B, “Vault’s File and Directory Structure,”](#) describes the directories and files installed with the Vault product.
- ◆ [Appendix C, “Vault Functional Design,”](#) describes the architectural and technical components of Vault.

Getting Help

VERITAS offers you a variety of support options.

Accessing the VERITAS Technical Support Web Site

The VERITAS Support Web site allows you to:

- ◆ obtain updated information about NetBackup Vault, including system requirements, supported platforms, and supported peripherals
- ◆ contact the VERITAS Technical Support staff and post questions to them
- ◆ get the latest patches, upgrades, and utilities
- ◆ view the NetBackup Vault Frequently Asked Questions (FAQ) page
- ◆ search the knowledge base for answers to technical support questions
- ◆ receive automatic notice of product updates
- ◆ find out about NetBackup Vault training
- ◆ read current white papers related to NetBackup Vault

The address for the VERITAS Technical Support Web site follows:

- ◆ <http://support.veritas.com>

Subscribing to VERITAS Email Notification Service

Subscribe to the VERITAS Email notification service to be informed of software alerts, newly published documentation, Beta programs, and other services.



Go to <http://support.veritas.com>. Select a product and click “E-mail Notifications” on the right side of the page. Your customer profile ensures you receive the latest VERITAS technical information pertaining to your specific interests.

Accessing VERITAS Telephone Support

Telephone support for NetBackup Vault is only available with a valid support contract. To contact VERITAS for technical support, dial the appropriate phone number listed on the Technical Support Guide included in the product box and have your product license information ready for quick navigation to the proper support group.

▼ To locate the telephone support directory on the VERITAS web site

1. Open <http://support.veritas.com> in your web browser.
2. Click the **Phone Support** icon. A page that contains VERITAS support numbers from around the world appears.

Accessing VERITAS E-mail Support

▼ To contact support using E-mail on the VERITAS web site

1. Open <http://support.veritas.com> in your web browser.
2. Click the **E-mail Support** icon. A brief electronic form will appear and prompt you to:
 - ◆ Select a language of your preference
 - ◆ Select a product and a platform
 - ◆ Associate your message to an existing technical support case
 - ◆ Provide additional contact and product information, and your message
3. Click **Send Message**.

Contacting VERITAS Licensing

For license information call 1-800-634-4747 option 3, fax 1-650-527-0952, or e-mail amercustomer@veritas.com.



NetBackup Vault Manuals

The companion document to this *NetBackup Vault System Administrator's Guide for Unix and Windows* is the *NetBackup Vault Operator's Guide for Unix and Windows*, which provides instructions about using NetBackup Vault for system operators.

The following documents provided related information:

- ◆ *VERITAS NetBackup Release Notes for UNIX and Windows*

Provides important information about NetBackup on UNIX- and Windows-based servers, such as the platforms and operating systems that are supported and operating notes that may not be in the NetBackup manuals or the online help.

- ◆ *VERITAS NetBackup Installation Guide for UNIX*

Explains how to install NetBackup software on UNIX-based platforms.

- ◆ *VERITAS NetBackup Installation Guide for Windows*

Explains how to install NetBackup software on Windows-based platforms. Also explains how to install PC client software, which includes UNIX systems and Mac OS 10.

If you have a UNIX server, refer to these documents:

- ◆ *VERITAS NetBackup System Administrator's Guide for UNIX, Volume I*

Explains how to configure and manage NetBackup on a UNIX server, including managing storage units, backup policies, catalogs and host properties.

- ◆ *VERITAS NetBackup System Administrator's Guide for UNIX, Volume II*

Explains additional NetBackup features such as notify scripts, enhanced authorization and authentication, and role-based security. The guide also discusses using NetBackup with AFS, Intelligent Disaster Recovery (IDR), and the BE Tape Reader.

- ◆ *VERITAS NetBackup Media Manager System Administrator's Guide for UNIX*

Explains how to configure and manage the storage devices and media on UNIX servers running NetBackup. Media Manager is part of NetBackup.

- ◆ *VERITAS NetBackup Troubleshooting Guide for UNIX and Windows*

Provides troubleshooting information for UNIX- and Windows-based NetBackup, including Media Manager.

- ◆ *VERITAS NetBackup Commands for UNIX*

Describes NetBackup and Media Manager commands and processes that can be run from a UNIX command line.

If you have a Windows server, refer to these documents:



- ◆ *VERITAS NetBackup System Administrator's Guide for Windows, Volume I*
Explains how to configure and manage NetBackup on a Windows server, including managing storage units, backup policies, catalogs and host properties.
- ◆ *VERITAS NetBackup System Administrator's Guide for Windows, Volume II*
Explains additional NetBackup features such as notify scripts, enhanced authorization and authentication, and role-based security. The guide also discusses using NetBackup with AFS, Intelligent Disaster Recovery (IDR), and the BE Tape Reader.
- ◆ *VERITAS NetBackup Media Manager System Administrator's Guide for Windows*
Explains how to configure and manage the storage devices and media on Windows servers running NetBackup. Media Manager is part of NetBackup.
- ◆ *VERITAS NetBackup Troubleshooting Guide for UNIX and Windows*
Provides troubleshooting information for UNIX- and Windows-based NetBackup, including Media Manager.
- ◆ *VERITAS NetBackup Commands for Windows*
Describes NetBackup commands and processes that can be executed from a Windows command prompt.

Glossary

If you encounter unfamiliar terminology, consult the NetBackup online glossary. The glossary contains terms and definitions for NetBackup and all additional NetBackup options and agents.

The NetBackup online glossary is included in the NetBackup help file.

▼ To access the NetBackup online glossary

1. In the NetBackup Administration Console, click **Help > Help Topics**.
2. Click the **Contents** tab.
3. Click **Glossary of NetBackup Terms**.

Use the scroll function to navigate through the glossary.



Accessibility Features

NetBackup contains features that make the user interface easier to use by people who are visually impaired and by people who have limited dexterity. Accessibility features include:

- ◆ Support for assistive technologies such as screen readers and voice input (Windows servers only)
- ◆ Support for keyboard (mouseless) navigation using accelerator keys and mnemonic keys

For more information, see the *NetBackup System Administrator's Guide for Windows, Volume I* or the *NetBackup System Administrator's Guide for UNIX, Volume I*.

Conventions

The following conventions apply throughout the documentation set.

Product-Specific Conventions

The following term is used in the NetBackup Vault 5.1 documentation to increase readability while maintaining technical accuracy.

- ◆ Microsoft Windows, Windows

Terms used to describe a specific product or operating system developed by Microsoft, Inc. Some examples you may encounter in NetBackup documentation are, Windows servers, Windows 2000, Windows Server 2003, Windows clients, Windows platforms, or Windows GUI.

When Windows or Windows servers is used in the documentation, it refers to all of the currently supported Windows operating systems. When a specific Windows product is identified in the documentation, only that particular product is valid in that instance.

For a complete list of Windows operating systems and platforms that NetBackup supports, refer to the *NetBackup Release Notes for UNIX and Windows* or go to the VERITAS support web site at <http://www.support.veritas.com>.

Typographical Conventions

Here are the typographical conventions used throughout the manuals:

Conventions

| Convention | Description |
|-----------------|---|
| GUI Font | Used to depict graphical user interface (GUI) objects, such as fields, listboxes, menu commands, and so on. For example: Enter your password in the Password field. |
| <i>Italics</i> | Used for placeholder text, book titles, new terms, or emphasis. Replace placeholder text with your specific text. For example: Replace <i>filename</i> with the name of your file. Do <i>not</i> use file names that contain spaces. This font is also used to highlight NetBackup server-specific or operating system-specific differences. For example: <i>This step is only applicable for NetBackup Enterprise Server.</i> |
| Code | Used to show what commands you need to type, to identify pathnames where files are located, and to distinguish system or application text that is displayed to you or that is part of a code example. |
| Key+Key | Used to show that you must hold down the first key while pressing the second key. For example: Ctrl+S means hold down the Ctrl key while you press S. |

You should use the appropriate conventions for your platform. For example, when specifying a path, use backslashes on Microsoft Windows and slashes on UNIX. Significant differences between the platforms are noted in the text.

Tips, notes, and cautions are used to emphasize information. The following samples describe when each is used.

Tip Used for nice-to-know information, like a shortcut.

Note Used for important information that you should know, but that shouldn't cause any damage to your data or your system if you choose to ignore it.

Caution Used for information that will prevent a problem. Ignore a caution at your own risk.

Command Usage

The following conventions are frequently used in the synopsis of command usage.



brackets []

The enclosed command line component is optional.

Vertical bar or pipe (|)

Separates optional arguments from which the user can choose. For example, when a command has the following format:

`command arg1|arg2`

In this example, the user can use either the *arg1* or *arg2* variable.

Navigating Multiple Menu Levels

When navigating multiple menu levels, a greater-than sign (>) is used to indicate a continued action.

The following example shows how the > is used to condense a series of menu selections into one step:

- ❖ Select **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.

The corresponding actions could be described in more steps as follows:

1. Click **Start** in the task bar.
2. Move your cursor to **Programs**.
3. Move your cursor to the right and highlight **VERITAS NetBackup**.
4. Move your cursor to the right. First highlight and then click **NetBackup Administration Console**.

Introduction to Vault

1

NetBackup Vault is an extension to NetBackup that automates selection and duplication of backup images and ejection of media for transfer to and from a separate, off-site storage facility. NetBackup Vault also generates reports to track the location and content of the media. Vault functionality does not have to be used only for disaster recovery; you can use Vault to manage data and media that you store off-site for regulatory archival purposes.

For more information, see the following:

- ◆ [“Vault Uses NetBackup Functions”](#) on page 1
- ◆ [“How to Access NetBackup Vault”](#) on page 2
- ◆ [“Vault Original or Duplicate Images?”](#) on page 2
- ◆ [“The Vault Process”](#) on page 2
- ◆ [“How Vault Uses Volume Groups and Pools”](#) on page 4
- ◆ [“NetBackup and Vault Configuration”](#) on page 5

Vault Uses NetBackup Functions

NetBackup Vault uses existing NetBackup functions for all operations, such as duplication of images, media control, reporting, and ejecting and injecting of tapes. Information from Vault is integrated with other NetBackup components and appears in the NetBackup Activity Monitor.

NetBackup Vault interacts with the following NetBackup services and catalogs:

- ◆ Media Manager manages robots and media.
- ◆ The NetBackup catalog and the Media Manager database record which images have been vaulted.
- ◆ The Media Manager database information determines when expired media can be returned to the robot for reuse.
- ◆ The Activity Monitor displays the status of the Vault job.



How to Access NetBackup Vault

NetBackup Vault is installed on a NetBackup master server. If you added the appropriate license key during the installation of Vault or by using the **Help > License Keys** option of the NetBackup Administration Console, **Vault Management** will be a node of the Administration Console. You can use the Administration Console to configure and manage Vault.

Alternatively, you can manage Vault by using the following methods:

- ◆ Menu-based user interfaces.
- ◆ Command line utilities.

For more information about installing NetBackup Vault on a UNIX server, see “[Installing NetBackup Vault on a UNIX System](#)” on page 8. For information about licensing NetBackup Vault on a Windows server, see “[Licensing Prerequisites for a Windows System](#)” on page 16.

Vault Original or Duplicate Images?

One of your most important choices is whether to send original or duplicate images off site. If you send original images off site, you do not have to duplicate images and therefore do not have to configure duplication. Vault distinguishes between original images and duplicate images as follows:

- ◆ *Original images* are created by NetBackup during a backup job, including all copies created concurrently by a backup policy. NetBackup can create up to four copies of an image concurrently during the backup process; all are considered originals.

By default, the first (or only) image made during a backup job is the primary backup. NetBackup restores from the primary backup, and Vault duplicates from the primary backup. Therefore, in most circumstances the primary backup should remain in the robot.

- ◆ *Duplicate images* are copies created by Vault. A Vault job reads the primary backup image and writes one or more duplicate images concurrently; the Vault job is separate from the NetBackup policy job.

The Vault Process

Vaulting is the process of sending backup images off site to a protected storage location. For more information about the specific steps in a Vault process, see the following, which briefly describes a basic Vault process:

- ◆ [Choose Backup Images](#)



- ◆ Duplicate Backup Images
- ◆ Backup the NetBackup Catalog
- ◆ Eject Media
- ◆ Generate Reports

A Vault job must select images (Choose Backups). The other steps are optional so you can separate the Vault tasks into separate jobs if desired, using different jobs to accomplish different tasks. For example, you can use one job to select and duplicate images daily, and another job to eject media and generate reports weekly.

Injecting returned media back into the robot is a manual operation. The Vault reports include the media that should be recalled from the off-site location and injected into the robot.

The term *vault* also refers both to a logical entity associated with a particular robot and to the off-site storage location of a set of tapes.

Choose Backup Images

The first step of the Vault process is choosing the backup images that are candidates to be transferred off site. This step, known as image selection, must be configured for every Vault job. Vault uses the criteria in a Vault *profile* (a set of rules for selecting images, duplicating images, and ejecting media) to determine which backup images are candidates to send off site.

If you create multiple original images concurrently during a backup job, Vault will choose one or more of the originals and send original images off site (depending on the profile rules). If you duplicate images, Vault will choose the primary backup image and use that as the source image for the duplication operation.

Duplicate Backup Images

The second step of the Vault process is duplicating the backup images that are candidates to be transferred off site. This step, known as image duplication, produces media that you can eject and transfer off site.

Image duplication is optional. If you create multiple backup images concurrently during a NetBackup policy job and send one or more of those off site, you do not have to duplicate images in Vault and therefore do not have to configure the duplication step. However, if you create only one set of backup images during a NetBackup policy job, you must configure Vault to duplicate images so you can maintain backup images in your library and also eject and send images off site.



Backup the NetBackup Catalog

The third step of the Vault process is backing up the NetBackup and Media Manager catalogs. The NetBackup and Media Manager catalogs consist of databases of information about the NetBackup configuration and any backups that have been performed. The information about backups includes records of the files and the media on which the files are stored, including information about media sent off-site. The catalogs also have information about the media and storage devices that are under the control of Media Manager.

Backing up the catalog is optional. However, vaulting a catalog backup with your backup data can help you recover from a disaster more efficiently. Vault creates its own catalog backup with up-to-date information; Vault does not duplicate the NetBackup catalog.

Eject Media

The fourth step of the Vault process is ejecting the media that you then transfer to secure storage, often at a separate facility. Media that are ejected are tracked by Vault reporting facilities and will be recalled from off-site storage for reuse after the images expire. Media can be ejected automatically by a scheduled Vault job or manually after the job has completed. Media can be ejected for each job individually or can be consolidated into a single eject operation for multiple Vault jobs.

Generate Reports

The fifth step of the Vault process is generating reports. Reports track the media managed by Vault. You and your off-site storage vendor can use the reports to determine which media should be moved between your site and the off-site storage location and the timing of the moves.

Reports can be generated as part of the Vault job or manually after the job is finished. Reports can be generated for each job individually or can be consolidated with a consolidated eject operation. Generating reports is optional.

How Vault Uses Volume Groups and Pools

Volume groups identify where volumes reside. They are used as a tracking mechanism by Vault to determine where a volume is located. Volumes in a *robotic volume group* reside in a robot. When a volume is ejected and sent to off-site storage, Vault moves it logically to an *off-site volume group*. (A *logical move* means to change the volume attributes to show the new location.) When a volume in off-site storage expires and is injected back into the robot, it is moved back into the robotic volume group.

Volume pools identify logical sets of volumes by usage. They are used by Vault to determine if a volume should be ejected. Volume pools for images to be transferred off site are known as *off-site volume pools*. Images that you want to send off site must be in an off-site volume pool. During a Vault job, Vault searches a robot for media that matches the selection criteria; if it belongs to an off-site volume pool, Vault ejects it.

NetBackup and Vault Configuration

Before you can begin to use Vault, you must first set up and configure NetBackup so that volume pools and policies are available to support Vault operations. To do so, see [“Preparing NetBackup for Vault”](#) on page 43.

After configuring NetBackup for use with Vault, you must configure Vault robots and profiles. [“Configuring Vault”](#) on page 49 provides instructions for configuring Vault.

You should also review the information in [“Best Practices”](#) on page 19; it can help you determine how to setup and configure Vault most effectively for your environment, resources, requirements, service level agreements, and so on.





Installing NetBackup Vault

This chapter outlines the steps required to install NetBackup Vault on both UNIX and Windows systems. The following topics are covered:

- ◆ [Supported Systems](#)
- ◆ [Supported Robots](#)
- ◆ [UNIX Systems](#)
- ◆ [Microsoft Windows Systems](#)
- ◆ [Upgrading from bpvault 3.4](#)

Supported Systems

NetBackup Vault runs on the same operating systems and versions and in the same clustering environments as NetBackup except as noted in the *NetBackup Release Notes*. NetBackup restrictions and limitations related to systems, clusters, and peripherals also apply to Vault. *Exception:* Vault does not support standalone drives.

For information about supported systems and supported upgrade paths, see the *NetBackup Release Notes*.

Supported Robots

NetBackup Vault supports all Media Manager-supported robot types except optical disk library (ODL) robots. Robots that do not have media access ports and barcode readers are supported by NetBackup Vault. However, for best performance and to avoid user errors when entering media IDs, VERITAS recommends that you use robots that have media access ports and barcode readers.

UNIX Systems

NetBackup Vault can be installed and uninstalled on a UNIX system.



Installation Prerequisites for a UNIX System

- ◆ A NetBackup master server must be installed on the UNIX system. Vault cannot be installed on a NetBackup media server or on a NetBackup client.
- ◆ You must have a valid NetBackup Vault license key.

Installing NetBackup Vault on a UNIX System

Use this procedure to do an initial installation on a UNIX system. NetBackup must be installed before you can install Vault, and you must install the same version of Vault as NetBackup. For instructions about installing NetBackup, see the *NetBackup Installation Guide for UNIX*.

If you are upgrading Vault, see [“Upgrading NetBackup Vault on a UNIX System”](#) on page 10.

If you are installing Vault in a cluster environment, you must install Vault on all systems in the cluster on which NetBackup master servers are installed.

For information about where the various Vault components are installed, see [“Vault’s File and Directory Structure”](#) on page 217.

▼ To install NetBackup Vault on a UNIX server

Note If you are installing Vault in a cluster environment, you must freeze the active node before you begin the installation process so that migrations do not occur during installation. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator’s Guide* for the cluster software you are running.

1. Log in as the root user on the system on which the NetBackup master server is installed.
2. Verify that a valid license key for Vault has been registered by executing the following command to list and add keys:

```
/usr/openv/netbackup/bin/admincmd/get_license_key
```

3. Insert the CD that contains the Vault software into the CD-ROM drive.
4. Change the working directory to the CD-ROM directory.

```
cd /cd_rom_directory
```

Where *cd_rom_directory* is the path to the directory where you can access the CD-ROM. On some platforms, you may need to mount this directory. For instructions on how to mount the directory, see the *NetBackup Installation Guide for UNIX*.

5. Enter the following command:

```
./install
```

6. Select the **NetBackup Add-On Product Software** option.

A menu of NetBackup product options is displayed.

7. Select the **NetBackup Vault** option.

8. Enter **q** to quit the menu.

9. When asked if the list is correct, answer **y**.

The installation process begins. When completed, the Installation Options menu appears.

10. Enter **q** to quit the Installation Options menu.

11. In a cluster environment environment, complete [step 1](#) through [step 10](#) for each node on which you are installing a NetBackup master server.

Note If you are installing Vault in a cluster environment, unfreeze the active node after the installation completes. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

12. Start the NetBackup Administration Console by entering the following command:

```
/usr/opensv/netbackup/bin/jnbSA&
```

To complete a new installation, you must configure appropriate NetBackup attributes for use by Vault and identify which NetBackup policies you want to use with Vault (or create new ones to use with Vault). Please read the following chapters so you will develop an understanding of how Vault works and how best to configure Vault for your operations. You should configure the e-mail address for notification of session status and enter alternate media server names, if appropriate. See "[Configuring Vault Properties](#)" on page 52.

In a cluster environment, you can configure Vault by using the NetBackup Administration Console connected through the NetBackup virtual server name, regardless of which cluster server is currently active.



Upgrading NetBackup Vault on a UNIX System

Use this procedure if you already have Vault installed and are upgrading to a newer version of Vault. Use one of the following procedures:

- ◆ To upgrade NetBackup Vault on a Solaris system
- ◆ To upgrade NetBackup Vault on a UNIX system other than Solaris

During an upgrade installation, you can choose to use the existing Vault license key.

▼ To upgrade NetBackup Vault on a Solaris system

Note If you are upgrading Vault in a cluster environment, you must freeze the active node before you begin the upgrade process so that migrations do not occur during the upgrade. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in as root to the system on which Vault is installed.
2. Execute the following command to uninstall NetBackup Vault:

```
pkgrm VRTSnbvlt
```

A message asking if you want to remove the package is displayed:

```
Do you want to remove this package?
```
3. Enter **y** to remove Vault.
4. If you are prompted with a message about super-user permissions and you are asked if you want to continue, enter **y**.
5. Messages that show the progress of the removal process are displayed until the following message appears:

```
Are you doing this pkgrm as a step in an upgrade process?  
[y,n,?,q]
```
6. Enter **y**.

The package removal process will remove Vault program components but not database and log files, so you will not lose your configuration during the upgrade.
7. Upgrade NetBackup by following the upgrade installation procedures for NetBackup in the *NetBackup Installation Guide for UNIX*.

8. Insert the CD that contains the Vault software into the CD-ROM drive.

9. Change the working directory to the CD-ROM directory.

```
cd /cd_rom_directory
```

Where *cd_rom_directory* is the path to the directory where you can access the CD-ROM. On some platforms, you may need to mount this directory. For instructions on how to mount the directory, see the *NetBackup Installation Guide for UNIX*.

10. Enter the following command:

```
./install
```

11. Select the **NetBackup Add-On Product Software** option.

A menu of NetBackup product options is displayed.

12. Select the **NetBackup Vault** option.

13. Enter **q** to quit the menu.

14. When asked if the list is correct, answer **y**.

The installation process begins. When completed, the Installation Options menu appears.

15. Enter **q** to quit the Installation Options menu.

16. In a cluster environment environment, complete [step 1](#) through [step 15](#) for each node on which Vault is installed.

Note If you are upgrading Vault in a cluster environment, unfreeze the active node after upgrading Vault. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

17. Start the NetBackup Administration Console by entering the following command:

```
/usr/opensv/netbackup/bin/jnbSA&
```



▼ To upgrade NetBackup Vault on a UNIX system other than Solaris

Note If you are upgrading Vault in a cluster environment, you must freeze the active node before you begin the upgrade process so that migrations do not occur during the upgrade. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in as root to the system on which Vault is installed.
2. Optionally, delete the following files and directories. Although not required, deleting these items can help ensure that all Vault files installed on the system are the same revision.

```
/usr/opensv/netbackup/bin/bpbrmvlt  
/usr/opensv/netbackup/bin/vlt*  
/usr/opensv/netbackup/bin/goodies/vlt*  
/usr/opensv/netbackup/help/vltadm
```

3. Upgrade NetBackup by following the upgrade installation procedures for NetBackup in the *NetBackup Installation Guide for UNIX*.
4. Insert the CD that contains the Vault software into the CD-ROM drive.
5. Change the working directory to the CD-ROM directory.

```
cd /cd_rom_directory
```

Where *cd_rom_directory* is the path to the directory where you can access the CD-ROM. On some platforms, you may need to mount this directory. For instructions on how to mount the directory, see the *NetBackup Installation Guide for UNIX*.

6. Enter the following command:

./install
7. Select the **NetBackup Add-On Product Software** option.
A menu of NetBackup product options is displayed.
8. Select the **NetBackup Vault** option.
9. Enter **q** to quit the menu.
10. When asked if the list is correct, answer **y**.

The installation process begins. When completed, the Installation Options menu appears.

11. Enter **q** to quit the Installation Options menu.
12. In a cluster environment environment, complete [step 1](#) through [step 11](#) for each node on which you are installing a NetBackup master server.

Note If you are upgrading Vault in a cluster environment, unfreeze the active node after upgrading Vault. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

13. Start the NetBackup Administration Console by entering the following command:

```
/usr/opensv/netbackup/bin/jnbSA&
```

In a cluster environment, you should run the NetBackup Administration Console on the primary node in the cluster when you perform Vault configuration.

Uninstalling NetBackup Vault from a UNIX System

The following instructions describe how to remove Vault permanently from a UNIX system without uninstalling NetBackup.

Before you uninstall NetBackup Vault, you should delete all Vault-specific items from NetBackup, such as volume pools, Vault policies, and so on. Although all Vault program components are removed when Vault is uninstalled, configuration items related to NetBackup are not.

All Vault components and configuration information is removed during this procedure. The procedure you should follow depends on whether the UNIX system runs the Solaris operating system or some other version of the UNIX operating system.

Use one of the following procedures:

- ◆ To remove NetBackup Vault from a Solaris System
- ◆ To remove NetBackup Vault from a UNIX system

▼ To remove NetBackup Vault from a Solaris system

Caution This process removes Vault completely, including the Vault database and log files.



Note If you are removing Vault in a cluster environment, you must freeze the active node before you begin removing Vault so that migrations do not occur during the removal process. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in as root to the system on which Vault is installed.
2. Execute the following command to uninstall NetBackup Vault:

```
pkgrm VRTSnbvlt
```

A message asking if you want to remove the package is displayed:

```
Do you want to remove this package?
```
3. Enter **y** to remove Vault.
4. If you are prompted with a message about super-user permissions and you are asked if you want to continue, enter **y**.

Messages that show the progress of the removal process are displayed until the following message appears:

```
Are you doing this pkgrm as a step in an upgrade process?  
[y,n,?,q]
```
5. To remove Vault completely and not upgrade to a newer version, enter **n**.

Messages that ask about removing Vault files will appear.
6. Enter **y** to each prompt to remove the Vault files.
7. In a cluster environment environment, complete [step 1](#) through [step 6](#) for each node on which Vault is installed.

Note In a cluster environment, unfreeze the active node after removing Vault from all systems. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

▼ To remove NetBackup Vault from a UNIX system other than Solaris

Caution This process removes Vault completely, including the Vault database and log files.

Note If you are removing Vault in a cluster environment, you must freeze the active node before you begin removing Vault so that migrations do not occur during the removal process. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in as root to the system on which Vault is installed.
2. Remove the following files and directories:

```
/usr/opensv/netbackup/bin/bpbrmvlt  
/usr/opensv/netbackup/bin/vlt*  
/usr/opensv/netbackup/bin/goodies/vlt*  
/usr/opensv/netbackup/help/vltadm  
/usr/opensv/netbackup/vault  
/usr/opensv/share/version_vault
```
3. Remove the Vault database directory, `/usr/opensv/netbackup/db/vault`.
4. Remove the Vault log directory, `/usr/opensv/netbackup/logs/vault`.
5. In a cluster environment environment, complete [step 1](#) through [step 4](#) for each node on which Vault is installed.

Note In a cluster environment, unfreeze the active node after removing Vault from all systems. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

Microsoft Windows Systems

NetBackup Vault is installed on a Windows system when NetBackup is installed; no separate installation procedure is required. However, to use Vault, you must enter a license key. Two types of license keys are available:

- ◆ A license key for the base NetBackup product and all NetBackup add-ons, such as Vault. If you have already installed NetBackup and entered the license key, Vault is ready for use. You do not need to enter any additional license keys.
- ◆ A license key specifically for the Vault option. If so, you will have to enter the Vault license key before you can use Vault (see [“Licensing NetBackup Vault on a Windows System”](#) on page 16).



For information about where the various Vault components are installed, see “[Vault’s File and Directory Structure](#)” on page 217.

Licensing Prerequisites for a Windows System

- ◆ A NetBackup master server must be installed and running on the Windows computer. Vault cannot be installed on a NetBackup media server or on a NetBackup client.
- ◆ You must have a valid NetBackup Vault license key.

Licensing NetBackup Vault on a Windows System

Although NetBackup Vault is installed during NetBackup installation, you cannot use Vault until you enter the appropriate license key.

Note If the license key for NetBackup Vault was included in the license key for the base NetBackup product, you do not have to perform this procedure.

▼ To add the Vault license key

1. From the NetBackup Administration console, choose **Help > License Keys**.
The NetBackup License Keys dialog box appears.
2. Click **New** to display the Add a new License Key dialog box.
3. Enter the NetBackup Vault license key.
4. Click **Add**.

The license key is displayed in the NetBackup License Keys dialog box.

5. Click **Close** to close the NetBackup License Keys dialog box.

To complete a new installation, you must configure appropriate NetBackup attributes for use by Vault and identify which NetBackup policies you want to use with Vault (or create new ones to use with Vault). Please read the following chapters so you will develop an understanding of how Vault works and how best to configure Vault for your operations. Be sure to configure the e-mail address for notification of sessions status and enter alternate media server names, if appropriate. See “[Configuring Vault Properties](#)” on page 52.

Upgrading NetBackup Vault on a Windows System

On Windows systems, NetBackup Vault is upgraded at the same time NetBackup is upgraded. Therefore, to upgrade NetBackup Vault, follow the upgrade installation procedures for NetBackup in the *NetBackup Installation Guide for Windows*.

Delicensing NetBackup Vault from a Windows System

NetBackup Vault can be removed from the base NetBackup product by deleting the license key from the list of current licenses. When the license key is removed, NetBackup Vault is no longer available for use. You can delicense Vault only if Vault was licensed with its own license key, separate from the base NetBackup product license key.

Caution If NetBackup Vault was included as part of the base NetBackup product, removing the license key will disable the entire NetBackup base product.

Before you delicense NetBackup Vault, you should remove all Vault-specific configuration items by using the NetBackup Administration Console to delete them. Deleting the Vault configuration ensures that NetBackup does not include anything that was configured for Vault, such as volume pools.

The license must be removed from the master server on which the Vault software was installed initially.

▼ To delicense NetBackup Vault

1. From the NetBackup Administration Console, choose **Help > License Keys**.
The NetBackup License Keys dialog box is displayed.
2. From the list of keys displayed, select the NetBackup Vault license key.

Caution If NetBackup Vault was included as part of the base product key, performing the following step will delete your base key and you will be unable to use NetBackup. If you do not want to delete the NetBackup license key, do not continue.

3. If you want to continue with the delicensing, click the **Delete** button.
The Vault license key is deleted from the Current Licenses dialog box, and **Vault Management** is no longer displayed in the NetBackup Administration Console tree.



Upgrading from bpvault 3.4

To upgrade from bpvault 3.4 to NetBackup Vault, first upgrade to a NetBackup Vault 4.5 release and then follow the upgrade path documented in the *NetBackup Release Notes*.



Best Practices

Vault can be configured to support how your computing or data center environment is set up and how it operates. A *best practice* recommendation that may provide benefit for one environment may not provide the same benefit for another. You should evaluate and implement any recommendations based on the benefit to your environment.

For more information, see the following:

- ◆ [“Vaulting Paradigm”](#) beginning on page 20
- ◆ [“Preferred Vaulting Strategies”](#) on page 20
- ◆ [“Ensure All Data is Vaulted”](#) on page 22
- ◆ [“Do Not Vault More Than You Need To”](#) on page 24
- ◆ [“Preparing for Efficient Recovery”](#) on page 26
- ◆ [“Defer Ejection”](#) on page 29
- ◆ [“Avoid Resource Contention During Duplication”](#) on page 29
- ◆ [“Avoid Sending Duplicates Over The Network”](#) on page 35
- ◆ [“Increase Duplication Throughput”](#) on page 37
- ◆ [“Maximize Drive Utilization During Duplication”](#) on page 39
- ◆ [“Use Scratch Volume Pools”](#) on page 40
- ◆ [“Ensure Report Integrity”](#) on page 40
- ◆ [“Generate the Lost Media Report Regularly”](#) on page 41



Vaulting Paradigm

How you set up and name your vaults and profiles depends on your operations. For example, if you maintain a customer database and a payroll database, you may choose to organize your vaults by data usage and your profiles by time periods, as follows:

| Vaults | Profiles |
|------------|----------|
| CustomerDB | Weekly |
| | Monthly |
| Payroll | Biweekly |
| | Monthly |
| | Yearly |

Alternatively, if your operations are organized geographically, you can set up your vaults by location and your profiles by data type, as follows:

| Vaults | Profiles |
|--------|------------|
| London | CustomerDB |
| | Payroll |
| Tokyo | CustomerDB |
| | Payroll |

Preferred Vaulting Strategies

Several strategies can help you reduce resource and time contention when you back up your data and when you vault your backup media. Although these strategies may not be advantageous for all situations, they can be very beneficial in many environments. VERITAS recommends that you use one of the following:

- ◆ Vault the original NetBackup backup media. Because you can produce multiple copies of images during a NetBackup policy job, fewer drives and less time may be required to create multiple original copies than duplicating media.
- ◆ Use disk staging. Send your backups to disk and then copy the data from disk to removable media. This strategy reduces the time that the backup process uses.



Vault Original Backups

For most situations, VERITAS recommends that you use a NetBackup policy to produce multiple original backup images and then use a Vault profile to eject and transfer one or more of the original images off site. In most situations, vaulting originals has the following advantages:

- ◆ Vaulting originals uses less drive time than duplicating backup images from the original tapes. For example, a backup job that creates two originals of a backup image uses two drives — two units of drive time. Conversely, a backup job that creates one original image uses one drive and a vault job that creates one duplicate of that original uses two drives — three units of drive time. Over time, duplicating backup images consumes more drive-time than writing multiple originals during a backup job.
- ◆ Vaulting originals avoids drive and media contention problems that may occur if backup and duplication operations occur at the same time. If you have enough time so that NetBackup backup jobs and Vault duplication jobs never occur at the same time, resource contention may not be an issue. However, if backup jobs and Vault jobs overlap, competing processes may try to use the same drive resources at the same time.
- ◆ Vaulting originals avoids configuring for duplication. In complex environments (such as with multiple media servers, multiple robots, or multiple retention period requirements), it can be difficult to configure the duplication steps of Vault profiles. It is possible to send large amounts of data over the network without careful configuration, although in storage area network (SAN) environments network traffic may not be an issue.

If you decide to create and vault original backups, see the following information before you configure Vault:

- ◆ [“Vaulting Original Backups in a 24x7 Environment”](#) on page 25
- ◆ [“Avoid Vaulting Partial Images”](#) on page 24

Use Disk Staging

Note This topic is about using a disk storage unit as a destination for backup images, *not* about using a disk staging storage unit.

Using disk staging for your backup jobs can help avoid resource contention between backup operations and Vault duplication operations. Disk staging is the process of first writing the backup images to a disk storage unit during a NetBackup policy job and then writing the images to removable media during a Vault job. The following are some of the advantages of disk staging over tape-to-tape duplication:

- ◆ Shortens backup time. Writing to disk is faster than writing to tape, so less time is needed for backing up.



- ◆ Minimizes tape drive usage. Sending the original copy to tape then duplicating to a second tape, requires one drive to make the first copy and two drives (a read drive and a write drive) to make the second copy.
- ◆ Reduces resource contention problems. Because fewer tape drives are needed and the backup time period is shorter, it can help avoid drive and media contention between backup jobs and Vault jobs that do duplication.
- ◆ Reduces expense. Because disk access is fast and disk space is less expensive than tape drives, it is often advantageous to send your backups to disk.

You can schedule your Vault sessions to duplicate the original disk backup images to two (or more) tapes: one on-site tape and one off-site tape. Also, you can configure the Vault profile to free up the disk space automatically for the next round of backups.

Ensure All Data is Vaulted

When you are setting up NetBackup Vault, you should be sure that you configure it to vault all of the information that you want transferred off-site.

Overlap the Time Window in the Profile

To ensure that all data is vaulted, overlap the time window in the profile.

A Vault profile uses a time range as one of the criteria for choosing the backup images to be vaulted. Vault does not duplicate or eject a backup image that already has a copy in the Off-site Volume Group; therefore, Vault will not process images that have already been vaulted by a previous session. Perhaps more importantly, backups that were not processed if a previous session failed will be processed when the profile runs again if the time window is long enough. Therefore, configure the time window to be the sum of the following:

- ◆ The longest expected downtime for a server or robot
- ◆ Twice the length of the frequency at which the profile runs

For example, if you have a profile that duplicates images daily and your longest expected downtime is three days, you should configure the time window to be at least five days. If a robot fails and requires three days to repair, the next time the profile runs it will select backup images that did not get vaulted during the three-day downtime. Configuring the window to be longer, such as seven days, provides even more resiliency. A longer time window forces Vault to search a larger list of images for vault candidates; although that will consume more processing time, the extra processing time may not be a problem in your environment because Vault is a batch process that does not demand immediate system response.

Consequences of Not Overlapping the Time Window: Missing Data

When a vault session gets delayed, some backup images may be missed if the time window does not allow Vault to select images from a wider time range. For example, suppose the time window for your daily profile extends from 1 day ago to 0 days ago. On Tuesday, the robot has mechanical problems and the Vault profile fails. Consequently, Monday night's backups are not vaulted. On Wednesday, the robot is fixed. When the next Vault session begins on Wednesday, it will only select backup images that were created during the previous 24 hours, so Monday night's backups are still not vaulted. If the profile's time window had spanned more than 1 day, the session would have picked up both Monday night's and Tuesday night's backups.

Resolve Multiple Names for a Single Server

Note Alternate media servers apply to NetBackup Enterprise Server only.

For every media server, you should add an entry on the **Alternate Media Server Names** tab of the Vault Properties dialog. At a minimum, there should be, for each media server, an entry that contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action will avoid a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized as a match for the criteria entered in the **Choose Backups** tab of the Profile dialog and may therefore not get vaulted.

If you have multiple NIC cards in your server, make sure that the server name or IP address associated with each NIC card is listed in the Alternate Media Server Names tab when you configure a profile.

For more information, see "[Alternate Media Server Names Tab](#)" on page 53.

Specify Robotic Volume Group When Configuring a Vault

Volumes are ejected only if they are in a robotic volume group *and* in one of the off-site volume pools specified on the profile **Eject** tab. Therefore, if you want a volume to be ejected, ensure that it is in a robotic volume group and in one of the off-site volume pools specified on the profile **Eject** tab.

Multiple Volume Groups (Multiple Robots)

A profile will only eject media from the robotic volume group of the vault to which the profile belongs, and a volume group cannot span robots (typically, a volume group identifies a specific robot). However, a profile can select images to duplicate that are in a different robot's volume group and in multiple volume groups, which is useful if you



have backup images on multiple robots and want to duplicate those images on media in a robot from which the media will be ejected. If you use this configuration, it must be configured with care as described in “[Alternative A: Dedicated Robot for Vault Processing](#)” on page 30.

If your profile does not do duplication, you do not have to specify a Source Volume Group on the **Choose Backups** tab; if you specify a Source Volume Group, it has no effect on images that are vaulted.

Do Not Vault More Than You Need To

When you are setting up NetBackup Vault, you should be careful that you do not select and transfer off-site more data than you need to.


Send Only the Intended Backups Off-site

When configuring your backup policies, do not assign backup images not intended to be moved off-site to volumes in an off-site volume pool. In some circumstances, Vault will eject a volume if it contains images not intended for off-site storage. For example, if volume ABC123 has three images from policy1 and three images from policy2, and policy1 is specified on the profile **Eject** tab, volume ABC123 will be ejected even though it contains images from policy 2.

Use different volume pools for backup images you want to keep on site and for backup images you want to send to the vault. If you use the same volume pool for both, you will vault the backup images that should remain on-site. Also, if you use the same volume pool for both, a deadlock situation may result if your Vault profile is duplicating images because it may attempt to read a backup image from the same tape to which it tries to write the image.

Avoid Vaulting Partial Images


Original backup tapes often begin and end with partial images, shown as follows:



If you eject and vault original backup media, that media may contain partial images. To avoid vaulting partial images, use one of the following methods:

- ◆ Stop backup activity long enough to run Vault.
- ◆ If backup jobs are running, use the Suspend Media for the Next Session option on the profile **Eject** tab to suspend all media on which backups were written within the last day and then vault only those backups that are older than one day. No more backup images will be written to that media, and that media will be ready to be ejected.

When Suspend Media for the Next Session is used, the following shows what occurs during NetBackup and Vault operations:



Only use the Suspend Media for the Next Session option if you eject original backup media and want to avoid vaulting partial images. You should carefully consider whether to use the Suspend Media for the Next Session option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again, prolonging the length of time required to suspend the media. Therefore, some partial images on vaulted media may be acceptable. If you use this option, it is possible that the original backup media vaulted will not be full.

This option does not suspend media that is in use, such as media to which NetBackup is writing backup images.

Note Vault only suspends media in off-site volume pools specified on the profile **Eject** tab.

Vaulting Original Backups in a 24x7 Environment

If you use Vault in an environment in which backups can occur 24 hours a day, seven days a week, a profile may try to eject media to which backups are being written. Because Vault cannot suspend media on which backups are currently being written, an error will occur



and partial images may be vaulted. The rest of the image will be vaulted the next time the profile runs if that tape is not busy and the choose backups time window is large enough to select the image again.

To avoid such problems when vaulting originals, choose backups that were created a day or more ago and suspend the media to prevent writing to the media. (This assumes that your backups will be complete by the time the Vault session runs.)

Preparing for Efficient Recovery

Preparing in advance can help you restore your data more quickly and easily. The following can help you prepare for recovery.

Vault NetBackup Catalogs

To recover backups most efficiently, you will need a current NetBackup catalog. Although it is possible to rebuild the catalog by importing all of your backup media manually, it is a time-consuming process. A current catalog backup is a critical component of an effective disaster recovery plan. Therefore, it is very important to vault a catalog backup frequently. Frequently depends on how often you vault media. For example, if you vault media three times a day, perhaps vaulting a catalog once a day is often enough.

Use only one vault to create a copy of the NetBackup catalogs. If you have a robot attached to the master server, use it for the Vault catalog backup because in most circumstances that master server creates the NetBackup catalog that remains on site. (See the discussion of NetBackup catalog backups in the *NetBackup System Administrator's Guide, Volume I*). The profile that creates the catalog backup must also eject the catalog backup media; because catalog backups are processed differently than data backups, Vault will only eject catalog media for the currently running profile.

In most circumstances, you do not need to retain vaulted catalog backups for the same length of time that you retain other vaulted backup media. Although you only need one catalog backup in your off-site vault, for extra protection, maintaining the three most recent catalog backups in your off-site vault is a good practice. (The Recovery Report for Vault lists only the three most recent catalog backups in the off-site vault regardless of how many actually reside in the vault.)

To retain only the three most recent catalog backups, use the **Catalog Backup** tab of the Profile dialog to configure the retention period so that only the three most recent catalog backups remain in your off-site vault.

Use Precise Naming Conventions for Volume Pools and Groups

Use names for the off-site volume pools that make it easy for others to recognize them as data to be transferred off site, and use names for the off-site volume groups that indicate the physical location of the data. Use “offsite” or “vault” in the names, and try to identify the purpose or data in the pools and groups. For example, Vaulted_Payroll, Vaulted_CustomerDB, 1_month_vault, and 7_year_vault are descriptive volume pool names. How you name pools and groups can help you (and others) organize and more easily identify media if you have to recover data after a disaster.

Match Volume Pools to Data Usage

Volumes are assigned to volume pools. To assist with recovery, create and use off-site volume pools that match your data usage (that is, the type of data). For example, if you maintain a customer database, you will probably want to restore all of your customer database at the same time if you recover from a disaster. All of your customer database backup data should be assigned to an off-site volume pool specifically for that data, and only backup images of the customer database should be assigned to that off-site volume pool.

This volume pool (for example, Vaulted_CustomerDB) can correspond to all profiles within a logical vault or to a single profile, depending on how your Vault environment is configured.

Designate a Primary Copy and Keep It On Site

The first (or only) original backup image is the primary backup. NetBackup always uses the primary backup for restore operations, and Vault always uses the primary backup for duplication operations (unless a disk image is available). Ensure that the primary backups remain on site in your robot. If the primary backup is off site, a user initiated restore operation will wait indefinitely for a mount of the off-site media.

If you create multiple original backups during a NetBackup policy job, do not assign the primary copy to an off-site volume pool (unless you intend to send the first original off site). If you assign the first original to an off-site volume pool, it will be ejected and will not be available for restore or duplication operations.

If your Vault profile duplicates media and you send the first original off site, configure Vault to designate one of the duplicate backups that remains on site as the primary copy.



Suspend Vaulted Media

Unexpired media that is recalled and injected back into the robot should be suspended so NetBackup will not write images to it. Suspending media before it is ejected also helps to prevent errors from ejecting media that is in use. Vault profile **Eject** tab options let you suspend the media that is ejected so you do not have to suspend it if it is recalled. You also can choose to suspend media before it is ejected so that partial images are not written to that media.

The following are the two suspend options available on the **Eject** tab:

| Option | Purpose |
|------------------------------------|--|
| Suspend this Session's Media | <p>To suspend media in the eject list for the current session. If you select Immediately, no more images will be written to the media. If you select At Time of Eject, images may be written to the media until the media are ejected; select At Time of Eject if you want the media sent off-site to be full.</p> <p>Because Suspend this Session's Media operates on media in the eject list, it does not use more CPU cycles selecting media to suspend.</p> |
| Suspend Media for the Next Session | <p>To prevent partial images from being written onto media that contains images to be vaulted. Use this option only if you vault original images and want to avoid vaulting partial images on backup media.</p> <p>You should carefully consider whether to use this option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again. Also, this option will not suspend media that is in use, such as media to which NetBackup is writing backup images. Therefore, some partial images on vaulted media may be acceptable.</p> <p>This option will suspend duplicate media created by Vault; however, the Suspend this Session's Media option is a better choice for suspending duplicate media because it does not use CPU cycles to select media to suspend.</p> <p>For information about how partial images can be written to media, see "Avoid Vaulting Partial Images" on page 24.</p> |

Note Vault only suspends media in off-site volume pools specified on the **Eject** tab.



Revault Unexpired Media

You should always revault media that was recalled from off-site storage and injected into the robot (for example, if you recalled a volume to use for a restore operation). If you do not eject the media and transfer it to your off-site vault location, it will not be available if media at your site are damaged.

Defer Ejection

Different robotic libraries have different capabilities. Some can perform multiple operations at the same time, others cannot. You can reduce the chances of resource contention and error conditions from busy robots by ejecting media during a dedicated time period when no other inject or eject operations occur, as follows:

- ◆ Defer the eject process for all sessions.
- ◆ Reserve a specific time period for Vault media ejection.
- ◆ Eject media by configuring a profile for ejection only or by ejecting media manually.
- ◆ Do not inject or eject other media while Vault is ejecting media.
- ◆ Do not inventory a robot while Vault is ejecting media.

To determine which robotic libraries supported by NetBackup can perform multiple operations simultaneously, see the *NetBackup Media Manager System Administrator's Guide*.

Avoid Resource Contention During Duplication

Note If you vault original backups, you do not have to use practices that avoid or reduce resource contention in Vault.

Following are the resources that you should consider when configuring duplication in Vault:

- ◆ Time (that is, when the operations occur)
- ◆ Media used
- ◆ Robots and drives
- ◆ Bandwidth

Various configurations can help you avoid resource contention. Also, a general principle that can help avoid resource contention is to wait until backups are completed before using Vault to duplicate or eject media.



When Two Processes Try to Use the Same Drive

Careful configuration of your environment can help avoid resource contention during Vault duplication, which can occur when two processes try to use the same drive at the same time. To avoid resource contention, follow the advice provided in one of the following alternatives:

These alternative configurations work well for multi-robot environments; they use available resources wisely and are unlikely to cause resource allocation problems.

Alternative A: Dedicated Robot for Vault Processing

Note Alternate read servers apply to NetBackup Enterprise Server only.

In a multi-robot environment, dedicate one robot strictly for vault processing. The media in this robot will contain only the duplicate backup copies that are to be ejected and sent to the off-site vault. This configuration works best in a storage area network (SAN) environment where all media servers have direct access to the vault robot because then the duplication step will not send data over the network.

There are two ways to achieve this configuration, as follows:

- ◆ Use a NetBackup policy to create multiple original backup images concurrently . Write the first backup image (the primary backup) to a storage unit that is not in the Vault robot. Write one of the other originals to the Vault robot and assign it to the off-site volume pool. Configure a Vault profile to eject all media in that vault's off-site volume pool. This configuration requires that all robots used be connected to the same NetBackup media server.
- ◆ Use Vault to duplicate images. Backup images will be duplicated from all other robots to the Vault robot. Use one of the following alternatives to configure Vault to perform duplication:
 - ◆ On the **Duplication** tab of the Profile dialog, do not select **Advanced Configuration** or **Alternate Read Server**. For each backup image, the media server that performed the backup will also perform the duplication. All media servers will send duplication data to the Destination Storage Unit media server. If the Destination Storage Unit media server is not the same as the media server that performed the backup, the data will be sent over the network.
 - ◆ On the **Duplication** tab of the Vault Profile dialog, specify the destination storage unit's media server as the **Alternate Read Server** but do not select **Advanced Configuration**. If the alternate read server also has access to all of the backup robots, no data will be sent over the network.



- ◆ On the **Choose Backups** tab of the Profile dialog, specify **All Media Servers** in the Media Servers list. On the **Duplication** tab, select **Advanced Configuration**, select **Alternate Read Server**, then create an entry for each media server in your environment. To avoid sending duplication data over the network, for each media server entry specify the destination storage unit's media server as the alternate read server; that server must have access to all the robots that hold the source images so they will be duplicated. Ensure that the total number of write drives specified in the Write Drives column for each entry does not exceed the number of drives in the Vault robot.

If you use this alternative, do not use Any Available storage unit in your backup policies unless only your Vault storage units are set to On Demand Only. Using Any Available for other storage units may cause images not intended for off-site storage to be written to the Vault robot. You can achieve the same behavior provided by Any Available storage unit by configuring your backup policies to use a storage unit group that includes all storage units except for the vault robot's (although if you use storage unit groups you cannot make multiple copies simultaneously).

Advantage

This configuration is most convenient for the operator, who can eject and inject tapes from only one robot, simplifying the tape rotation process.

Disadvantage

In a complex environment, this alternative can be difficult to configure if you want to avoid sending duplication data over the network.

Alternative B: Each Robot as a Vault Robot

Note Alternate read servers apply to NetBackup Enterprise Server only.

In a multi-robot environment, configure each backup robot to be a Vault robot. Each robot will duplicate and/or eject only backup images that were originally written to it. You can do so in several ways, as follows:

- ◆ Use a NetBackup policy to create multiple original backups, assigning the copy to be vaulted to an off-site volume pool in any of the robots. For each robot, configure one vault and one profile that ejects the backups that were assigned to the off-site volume pool in that robot. Only backups on media in the off-site volume pools specified on the **Eject** tab and that meet the rest of the criteria specified in the profile will be ejected.



- ◆ Use Vault to duplicate images. On the **Choose Backups** tab of the Profile dialog, specify the robot to which the profile belongs in the **Source Volume Group** field. This will limit the profile so that it will duplicate only backup images that have their primary copy on media in this robot. Specify half of the available drives in the robot as read drives so that an equal number of read and write drives are available. Configure one such vault and profile for each robot.

To avoid sending duplication data over the network, specify the media server of the destination storage unit as the **Alternate Read Server**.

Advantages

These methods work well with backup policies that use Any Available storage unit. Using Vault to duplicate images also works well with storage unit groups if you make one copy only.

This configuration avoids resource contention when one profile attempts to duplicate images in multiple robots.

Note The destination storage unit must have at least two drives if that robot will be used for both read and write functions.

Alternative C: One Robot as Both a Backup and Vault Robot

In a multi-robot environment, configure all of the robots as backup robots and configure one of the backup robots as a Vault robot also. (One of the robots functions as both a backup robot and a Vault robot.) Configure one vault for the Vault robot, and in that vault configure one profile for each of the backup robots. In each profile, specify the backup robot in the **Source Volume Group** field of the **Choose Backups** tab and specify a destination storage unit that is in the Vault robot.

For example, if you have three robots that each have four drives, configure the three profiles as follows:

- ◆ In the profile for robot one (a backup robot only), specify the volume group in robot one as the Source Volume Group, specify four read drives, and specify a destination storage unit in robot three (robot three is the Vault robot). Images in robot one are read by four drives and written to four drives in robot three. Four duplication jobs run simultaneously.
- ◆ In the profile for robot two (a backup robot only), specify the volume group in robot two as the Source Volume Group, specify four read drives, and specify a destination storage unit in robot three. Images in robot two are read by four drives and written to four drives in robot three. Four duplication jobs run simultaneously.

- ◆ In the profile for robot three (a backup and Vault robot), specify the volume group in robot three as the Source Volume Group, specify two read drives, and specify a destination storage unit in robot three. Images in robot three are read by two drives and written to two drives. Two duplication jobs run simultaneously.

All images are duplicated to robot three and ejected from robot three.

Advantages

This method works well with backup policies that use Any Available storage unit. Using Any Available storage unit in your backup policies sends backup images to media in any storage unit available, and this configuration selects backup images on all the robots and duplicates them to the Vault robot.

Note The destination robot must have at least two drives if that robot will be used for both read and write functions.

When the Read Drive Is Not in the Vault Robot

The read drive does not have to be in the vault's robot. For configurations that include multiple media servers and multiple robots, we recommend that you seek advice from VERITAS Enterprise Consulting Services.

Sharing Resources with Backup Jobs

NetBackup and Vault use the same resources, specifically media servers and tape drives. Any Available storage unit can send some original backup images to the vault robot. Subsequently, when Vault tries to duplicate those images, it requires a read drive and a write drive in the vault robot. If not enough drives are available, a deadlock condition can occur. To prevent a potential conflict, ensure that backup jobs and Vault jobs are not scheduled at the same time. Vault contains load-balancing logic that requires one Vault session to use all configured read and write drives until the last of the tapes is duplicated. Although this is an efficient use of resources, NetBackup and Vault can try to use the same drive if backup jobs and Vault jobs are scheduled concurrently.

VERITAS recommends that you review a snapshot of the images you want to duplicate before you run the Vault job, which will show you where the images are located and what kind of resources will be required to duplicate them. You can create a snapshot using the preview method discussed in [“Previewing a Vault Session”](#) on page 100.



Load Balancing

If it is feasible, VERITAS strongly recommends that you create multiple original backup images concurrently in your backup policies to create both the on-site copy and the copy that will be sent to the vault rather than using Vault duplication. Avoiding the duplication step in a Vault profile avoids all resource contention issues and can significantly simplify the vaulting process.

If you cannot Vault originals, several strategies can help you balance the load on your computing environment.

Profiles for Both Originals and Duplicates

Vault can eject both original backups and duplicate images, so you can spread the load between backup jobs and Vault duplication jobs. For example, if your backup window is too small to create multiple simultaneous copies of all backups, you can create multiple copies of some of the backups and only one copy of the other backups and then configure a Vault profile to duplicate from the single original backups and eject both the original images and the duplicate images. For example:

- ◆ NetBackup policy A creates multiple original copies and assigns one of the copies to an off-site volume pool.
- ◆ NetBackup policy B creates one copy and assigns it to an on-site volume pool.
- ◆ Your Vault profile is configured to copy backup images and assign the duplicate images to an off-site volume pool.

When you run that Vault profile, Vault copies backup images from NetBackup policy B only and does not duplicate images from policy A because an original already exists in the off-site volume pool. If you have configured the profile for eject, it will eject both the copy of the original media from policy A and the duplicate media from policy B.

If Your Vault Vendor Does Not Pick Up Media Every Day

You can use Vault to duplicate backup images daily and eject volumes weekly. Duplication occurs every day rather than one day only, spreading the workload evenly throughout the week. The media remains in the robot until it is due to be collected by the vault vendor. For example, if the vault vendor picks up the media every Friday, you can do the following:

- ◆ Configure a Vault profile to do duplication only, and configure a vault policy to run this profile every day of the week.
- ◆ Configure a second Vault profile to do the catalog backup and eject steps. This profile should use the same image selection criteria as the profile that duplicates images. Configure a Vault policy to run this profile before the vault vendor arrives on Friday.

This method for duplicating and ejecting media provides the added benefit of consolidated reports that are not organized by session.

Specifying Different Volume Pools for Source and Destination

You should never configure a profile for duplication so that the source volume and the destination volume are in the same volume pool. This will result in deadlock when NetBackup chooses the same tape as the source and the destination of the duplication operation. (This is a NetBackup limitation.)

Avoid Sending Duplicates Over The Network

Sending duplicate images over the network is not a problem if there is sufficient bandwidth, but even a fiber optic storage area network (SAN) has only enough bandwidth for two or three duplication jobs at a time.

Following are some strategies you can use to avoid sending data over the network:

Create Originals Concurrently

One way to avoid sending data over the network with your Vault job is to create multiple original backup images concurrently during your scheduled backup jobs. This avoids the need for your Vault session to do duplication. In this scenario, Vault need only eject the backup tapes. Vault takes no significant resource time, except for the Catalog Backup step. (Catalog Backup is necessary to capture the changed volume database information for each vaulted tape.)

Suppose you want the on-site copy of your backups to go to one robot, and the off-site copy to go to another robot. If you create multiple backup images concurrently, all destination storage units must be on the same media server. Therefore, your media server will need a storage unit on both robots (one storage unit for your on-site copy and one for the off-site copy).

Use Alternate Read Server

Note Alternate read servers apply to NetBackup Enterprise Server only.

An alternate read server is a server used to read a backup image originally written by a different media server. You can avoid sending data over the network during duplication by specifying an alternate read server if the alternate read server is:


- ◆ Connected to the robot that has the original backups (the source volumes).



- ◆ Connected to the robot that contains the destination storage units.

Note If the destination storage unit is not connected to the alternate read server, you will send data over the network.

For example, in the diagram below, non-disk images written by media servers A, B, C, and D will be read by the alternate read server.




Use Advanced Duplication Configuration

Note More than one media server applies to NetBackup Enterprise Server only.

If each media server has access to at least one unique drive in the destination robot, you can use advanced duplication to process each media server independently and concurrently. (Note: all media from a single profile are ejected from the same robot.) You can do the same thing by configuring a separate profile for each media server rather than using advanced duplication configuration. However, multiple profiles within a single vault must run consecutively, so this may not allow you sufficient bandwidth.

In the following diagram, no alternate read server is used and each media server reads and duplicates its own backup images.



Take Care When Specifying All Media Servers

If you specify All Media Servers on the **Choose Backups** tab of a profile and also use Advanced Configuration on the **Duplication** tab, create an entry for each media server on the **Duplication** tab advanced configuration view.

If you list more media servers on the **Choose Backups** tab than on the **Duplication** tab, Vault assigns the images written by media servers not listed in the advanced view to the first media server that finishes its duplication job. If the first available media server is across the network, a lot of data would be sent over the network.

Another possible, though less problematic, consequence is that backup images from the media servers not configured for duplication may be duplicated by a different media server each time the profile is run.

Increase Duplication Throughput


Adding drives will enable Vault to run multiple duplicate sessions concurrently. For each write drive, a separate duplication job (bpduplicate) will be started. The following provides information about multiple drive environments:



Configuring for Multiple-Drives: Basics

In a basic multiple-drive configuration, there are an equal number of read and write drives, one master server, and one media server. The storage units are attached to the host on which the media server resides. A duplication process runs for each read/write drive pair. If the master server and media server reside on different hosts, media duplication data will be transferred over a network.

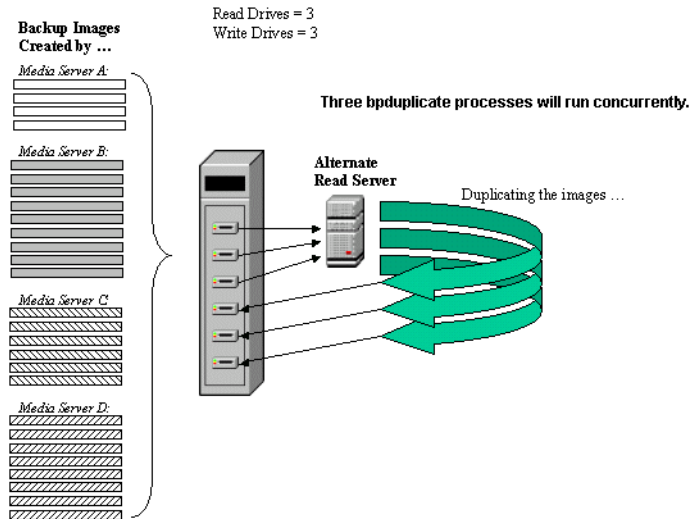
Note Only NetBackup Enterprise Server allows a master server and media server to reside on different hosts.



Multiple-Drive Scenario: Does Not Send Data Over Network

In a multiple-drive configuration that does not send data over the network, the configuration will have an equal number of read and write drives, one master server, and multiple media servers. A separate duplication process runs for each read/write drive pair during a duplication operation. If you designate an alternate read server (media server A) for reading the images to duplicate and if the destination storage unit also resides on the alternate read server (media server A), no data will be sent over the network.

Note Alternate read servers apply to NetBackup Enterprise Server only.



Maximize Drive Utilization During Duplication

To maximize drive utilization, VERITAS recommends that you do your duplication with as few Vault jobs as possible.

The more profiles you use, the less efficient the duplication process becomes. Drives will be idle between the duplication steps of consecutive Vault jobs while Vault is doing all of its other processing (selecting images, backing up the catalog, generating reports, and so on). It is much more efficient to use as few Vault profiles as possible for duplication. Therefore, if you can configure one Vault profile to duplicate all of your data, you will reduce idle time and get the maximum utilization of your drives.

In Vault 5.0 and later, you can configure one Vault profile to create off-site copies with multiple, different retentions. By doing this, a single Vault profile can do all of your duplication, which keeps your drives spinning from the time of the first image to the last, with no pause. For more information about multiple retention mappings, see [“Assigning Multiple Retentions with One Profile”](#) on page 121.



Use Scratch Volume Pools

A scratch pool is an optional volume pool that can be configured within VERITAS NetBackup. If a scratch pool is configured, Media Manager moves volumes from the scratch pool to other pools that do not have volumes available, including Vault pools. Expired volumes are returned to the scratch pool automatically. Because Media Manager allocates volumes from the scratch pool to the volume pools, a scratch pool helps ensure that volumes are available when needed.

You can set up a scratch pool in two ways, as follows:

- ◆ Create a scratch pool and add all your volumes to it. Create all the other volume pools but do not allocate any volumes to them. Media Manager will then move volumes from the scratch pool to the other volume pools as needed and return the expired volumes to the scratch pool.
- ◆ Create your volume pools and allocate volumes to them. Create a scratch pool and allocate volumes to it. Media Manager moves volumes between the scratch pool and the other volume pools as needed and returns the expired volumes to the scratch pool. This method may be the best option if you decide to add a scratch pool to an existing NetBackup configuration.

The scratch pool feature can affect reports for media coming on site. If you use a scratch pool, the Picking List for Vault, the Offsite Inventory, and the All Media Inventory reports may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even though the reports may be for a specific Vault profile or session.

Note If you use a scratch pool, Vault will allocate volumes from the scratch pool to the catalog volume pool if needed. However, Vault will not return expired catalog volumes back to the scratch pool because NetBackup never overwrites catalog media, even expired Vault catalog backup media. Vault returns expired catalog volumes back to the catalog volume pool.

For information about configuring NetBackup to use scratch pools, see the *NetBackup Media Manager System Administrator's Guide*.

Ensure Report Integrity

You should determine whether you want your Vault reports to group media by robot, by vault, or by profile. Your decision affects how you use your volume groups and volume pools.

Vault searches the off-site volume group for the media to include in the reports. It also uses the off-site volume pools for the same purpose. Therefore, you can use either the off-site volume group or the off-site volume pool(s) to organize media for each robot, vault, or profile.



Organizing Reports by Robot

To ensure that reports are organized by robot, all the vaults within each robot should use the same off-site volume group (that is, each robot has its own off-site volume group). This organizes reports by robot and maximizes the reuse of tapes. Media from one robot will not appear on the reports for another robot.

Reports will not seem consistent for an individual logical vault, but this strategy will maximize the frequency with which tapes are returned for reuse. Every time the Picking List for Vault report is generated for any profile within any vault for the robot, tapes from all profiles and logical vaults for that robot could be recalled for reuse (depending on how profiles share off-site volume pools).

Organizing Reports by Vault

To ensure that the Vault reports include media for each vault, specify a separate off-site volume group for each vault within a robot (that is, each vault has its own off-site volume group) and a common off-site volume pool for all profiles within each vault (that is, all profiles in the vault use the same off-site volume pool). Doing so ensures that each report contains media from one vault.

Organizing Reports by Profile

If you want the reports to include only media for a single profile, use a separate off-site volume pool for each profile.

Consequences of Sharing an Off-site Volume Group Across Multiple Robots

If profiles from multiple robots share both an off-site volume group and one or more off-site volume pools, your vault vendor will return a group of tapes (for a single Picking List for Vault report) that were ejected from multiple robots. The operator will need to identify which tapes should be injected into each of the robots. If mistakes are made identifying and injecting tapes, you can inject the incorrect media and possibly the incorrect number of media into your robots.

Generate the Lost Media Report Regularly

You should generate the Lost Media Report regularly so you can recall media that has not been returned from the off-site vault vendor but should have been returned. Media can get stranded at the off-site vault for various reasons:



- ◆ Frozen backup tapes never expire. A backup tape that does not expire will not appear on the Picking List for Vault and will not be recalled from the vault.
- ◆ A backup tape appears on the Picking List for Vault only once. If a tape from that report is missed and is not returned to the robot, it will never again be listed for recall.
- ◆ You change off-site volume group or pool names. If you change names, media may be stranded off-site because the Picking List for Vault is based on off-site volume pools and off-site volume groups, and media associated with the old names will not be listed.

How often you generate the Lost Media Report depends on your operations. Weekly or monthly may be often enough.



Preparing NetBackup for Vault

Before you can configure NetBackup Vault and begin using the Vault features, you must first configure NetBackup for your backup operations. You also must configure NetBackup so that it includes the appropriate volume pools, volume groups, and policies for use with Vault. Information in this section describes how to configure the NetBackup components that Vault requires; these components are *not* configured by using the Vault Management node of the Administration Console window.

Before you configure NetBackup for use with Vault, you should review the information in “[Best Practices](#)” on page 19. It can help you determine how to set up and configure policies, volume pools, and volume groups.

You should be familiar with basic NetBackup concepts, such as volume pools and groups, policies, and storage units. For more information about them, including how to configure them in NetBackup, see the *NetBackup System Administrator's Guide, Volume I* and the *Media Manager System Administrator's Guide*.

Volume Pools

A volume pool identifies a logical set of volumes by usage. Associating volumes with a volume pool protects them from access by unauthorized users, groups, or applications. Vault requires specific, dedicated volume pools. You can create new volume pools, or you can use volume pools that already exist in NetBackup if the pools are for Vault's exclusive use.

In Vault, volume pools for data to be transferred off site are known as *off-site volume pools*. During a Vault job, Vault searches for media that contains images that were created during the specified time period; if found and if it belongs to the specified off-site volume pool, Vault selects that media as a candidate for ejection.

How many volume pools you configure and how you name them depends on how your backup operations are configured and on how you want Vault to function. Following are basic volume pool usages:

- ◆ A volume pool for backup media that remains on site. If you are adding Vault to an existing NetBackup configuration, you probably already have volume pools that you can use for media that remains in the robot.



If you create multiple original backups in a NetBackup policy job, you should assign the first original (the primary backup) to the volume pool for media that remains on site. NetBackup restores from the primary backup, and Vault duplicates from the primary backup. Because both NetBackup and Vault use the primary copy, in most circumstances the primary copy should be the copy that remains in the robot.

- ◆ A volume pool for backup media transferred off site (an off-site volume pool). Vault ejects media from off-site volume pools, so data that you want to transfer off site should be assigned to an off-site volume pool. Choose names that clearly denote the volume pool as specifically for Vault and off-site storage (such as `Offsite_Backups`).

You can use this volume pool either for original backup images created as part of a NetBackup policy job or for duplicate images created by a Vault job.

- ◆ A volume pool for catalog media transferred off site (an off-site catalog volume pool). Vault requires a dedicated volume pool for catalog backups; you can assign catalog backups for all vaults to the same off-site catalog volume pool. Choose an easily identified name for the catalog volume pool (such as `Offsite_Catalog_Backup`). If you do not create and transfer off site the NetBackup catalogs, you do not need an off-site catalog volume pool.

After a tape is assigned to an off-site volume pool it remains in that pool and will be used for rotation within that same vault (unless a scratch pool exists, in which case it will be returned to the scratch pool).

When you name off-site volume pools for Vault, you should choose names that denote that the volume pools are specifically for off-site storage. VERITAS recommends that you use terms such as “offsite” or “vault” in the volume pool names to identify their purpose.

Do not use the NetBackup volume pool for Vault media. Because the NetBackup volume pool is the default volume pool, if you use it for Vault you will probably send more data off-site than you want to.

You should also record the following information about the volume pools so you can use it when you configure Vault:

- ◆ The names of the volume pool or pools
- ◆ The type and density of media in each volume pool

Volume pools are configured in the **Media and Device Management > Media** node of the NetBackup Administration Console. For information about how to create and configure volume pools, see the *Media Manager System Administrator's Guide*.

Best Practices

- ◆ [“Use Precise Naming Conventions for Volume Pools and Groups”](#) on page 27
- ◆ [“Match Volume Pools to Data Usage”](#) on page 27

Volume Groups

Volume groups identify a set of volumes that reside at the same physical location. Volume groups allow Vault to move volumes logically between a robot library and off-site storage (a *logical move* means to change the volume attributes to show the new location). Physical movement is accomplished when Vault ejects the media and an operator removes it and transfers it off-site.

Vault uses *robotic volume groups* and *off-site volume groups* to specify and determine the different locations at which volumes reside. A robotic volume group is a set of media that is located in a robot. An off-site volume group is a group of media that is stored at a different location; an off-site volume group is also known as a standalone volume group because it is not under the control of the robot.

When you create a vault (a *vault* is a logical entity associated with a specific robot), you must select both the robotic volume group and the off-site volume group. During a Vault job, Vault searches the robotic volume group for media that matches a profile's criteria; if media are found, Vault ejects that media and then moves it logically to the off-site volume group.

Usually, NetBackup creates a volume group name when volumes are added to a robot or when a robot that contains volumes is added to NetBackup, so you may not have to create a volume group to indicate that volumes are in a robot. NetBackup generates a name using the robot number and type. For example, if the robot is a TS8 and has a robot number of 50, the group name will be 00_050_TS8. A robotic library can contain volumes from more than one volume group, so a robot can have more than one robotic volume group name associated with it.

Conversely, you must name the off-site volume groups. When you create a vault, you select the robotic volume group from a list of volume groups for that robot. You also enter the name of an off-site volume group into which the volumes are moved after they are ejected from the robot. If the off-site volume group does not exist, it will be created during the vault session. The name of the off-site volume group should describe the data, the vault vendor, the vault location, or a combination thereof so you can easily identify the volume group. VERITAS recommends that you use terms such as "offsite" or "vault" in the volume group names to clearly identify their purpose.

For information about volume groups, see the *Media Manager System Administrator's Guide*.

Best Practices

- ◆ ["Use Precise Naming Conventions for Volume Pools and Groups"](#) on page 27



Vault Policies

A Vault policy is a NetBackup policy that is configured to run Vault jobs; a Vault policy does not back up client systems. The policy includes the schedule for when the Vault session should run (day or date and time window) and the command to run a Vault profile.

How many policies you use for Vault depends on how you conduct operations:

- ◆ A Vault policy can run a profile that ejects media that contains original images created during a backup job. If you create multiple original backup images concurrently, you can assign one or more of the original images to an off-site volume pool, and a separate Vault policy can run a Vault job that ejects the media on which those images are stored.
- ◆ A Vault policy can run a profile that selects images, duplicates those images, and ejects the media on which those images are stored. That policy can perform both operations daily or at some other interval that meets your requirements. If your vault vendor arrives daily to pick up media or you remove the off-site media from your robot immediately, you may need only one policy for that vault.
- ◆ One Vault policy can run a profile that duplicates images, and another policy can run a profile that ejects media. For example, if you create backup media daily and transfer it off site weekly, you can use one policy to create the backups daily and another policy to eject media weekly. If your vault vendor transfers your media weekly, you may prefer to eject media on the day the vault vendor arrives.
- ◆ If you transfer a catalog backup off site and the catalog backup requires more than one media volume (such as a tape), you will need a separate Vault policy for the catalog backup.

Policy Configuration Information

Collect and record the following information for each new schedule/policy pair you create or existing pair you consider for off-site rotation. The information you record will be used when you configure Vault and can help you determine if an existing policy can be used to create backup media that Vault can select for ejection.

Policy Configuration Information

| Property | Description |
|----------------|--|
| Policy Names | The names of all policies used for off-site rotation. To get information about existing policies, you can use the <code>bpp1list</code> command. |
| Schedule Names | The names of the schedule or schedules associated with each policy. |

Policy Configuration Information (continued)

| Property | Description |
|---------------------------------------|---|
| Off-site Original, Duplicate, or Both | Record whether the policy selects original backup media, creates duplicate backup media, or both. |
| Storage Unit | The storage units for each policy. |
| Retention Period | The retention period for each schedule so that you will have an idea of when to expect the media to return from off-site. |

Creating a Vault Policy

Setting up a Vault policy differs from setting up a regular policy in NetBackup. First, you must specify Vault as the policy type. Second, you do not specify clients for Vault policies. Third, rather than specifying files to backup on the **Backup Selections** tab, you specify one of two Vault commands to execute, as follows:

- ◆ Use the `vltrun` command to run a Vault session. You specify the robot, vault name, and profile for the job. The `vltrun` command accomplishes all the steps necessary to vault media. If the profile is configured for immediate eject, media are ejected and reports are generated. If the vault profile name is unique, use the following format:

```
vltrun profile_name
```

If the vault profile name is not unique, use the following format:

```
vltrun robot_number/vault_name/profile_name
```

- ◆ Use the `vlteject` command to eject media and/or generate reports for sessions that have been completed already and for which media have not been ejected. The `vlteject` command can process the pending ejects and/or reports for all sessions, for a specific robot, for a specific vault, or for a specific profile. For example:

```
vlteject -vault vault_name -eject -report
```

For more information about ejecting media, see [“Ejecting Media by Using a Vault Policy”](#) on page 110. For more information about the `vlteject` and `vltrun` commands, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual. For more information about creating NetBackup policies, see the *NetBackup System Administrator's Guide, Volume I*.

Note If you create a vault policy by copying a regular NetBackup policy that has a client list configured, delete all the clients in the client list before you run the policy. If you do not, Vault will create one vault job for every client in the list even though the



client is not used by the Vault job. The first vault job will operate as a normal vault session; the rest will terminate with a status 275 (a session is already running for this vault).

▼ To create a Vault policy

1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Click the **New Policy** button.
The Add a New Policy dialog appears.
3. Enter a unique name for the new policy in the Add a New Policy dialog.
4. Click **OK**.
The Add New Policy dialog appears.
5. On the **Attributes** tab, select **Vault** as the policy type.
6. On the **Schedules** tab, click **New** to create a new schedule.
The type of backup defaults to **Automatic Vault**.
7. Complete the schedule.
8. Bypass the **Client** tab (clients are not specified for Vault jobs).
9. On the **Backup Selections** tab, enter the appropriate Vault command for the policy.
10. Click **OK**.

Configuring Vault

When you configure Vault, you configure robots, vaults, and profiles.

Before you use Vault, you must configure NetBackup volume pools, volume groups, and policies for use with NetBackup Vault. To do so, see [“Preparing NetBackup for Vault”](#) on page 43.

You should also review the information in [“Best Practices”](#) on page 19; it can help you determine how to set up and configure Vault.

To configure Vault, see the following:

- ◆ [“Information Required to Configure Vault”](#) on page 49
- ◆ [“Configuring Vault Properties”](#) on page 52
- ◆ [“Configuring Robots for Vault”](#) on page 56
- ◆ [“Creating a Vault”](#) on page 57
- ◆ [“Creating a Profile”](#) on page 61
- ◆ [“Configuring a Profile”](#) on page 62

Information Required to Configure Vault

Information about the general configuration for NetBackup is required so you can set up and use NetBackup Vault. Collect and record the appropriate information about the following so that it is available when you begin to configure Vault.



Master Server, Media Servers, and Storage Units

Collect the following information about master servers, media servers, and robotic devices, which are used in various configuration options in Vault.

Server and Storage Unit Information

| Property | Description |
|---|---|
| Master Server Host Name | The name of the host server on which the NetBackup master server and Vault are installed. |
| Operating System Level of Master Server | The release of the operating system on the system on which the NetBackup master server is installed. |
| Number of Media Servers | The number of media servers associated with the master server. |
| Media Server Name | The name of each media server that controls the drives you want to use for the vault process. This server should also be bound to a storage unit within the NetBackup configuration. For NetBackup, all drives (of a given media type) that are attached to a server are defined as one storage unit, which is the recommended configuration for NetBackup. For every media server, configure alternate media server names; for more information, see " Alternate Media Server Names Tab " on page 53. |
| Operating System Level of Media Servers | The release of the operating system on the host machines on which the NetBackup media server or servers are installed. |
| Types of Robotic Devices | The robotic devices associated with each media server. Use the appropriate NetBackup terminology to identify the devices (for example, TLD, ACS, TL8) or specify the actual hardware manufacturer and model names for each device. |
| Storage Unit Name | The NetBackup storage units that are associated with each media server. You can use the <code>bpstulist -U</code> command to generate a list of existing storage units. Consider how many drives in each storage unit you want to use for vault sessions. You may choose to keep some drives available for restores or backups while duplication is running. |
| Number of Drives | The number of drives in each storage unit. Duplication requires drives in pairs: one to read and one to write. |



Robot Information

Collect the following information for each robot. Although the following information is not required to configure a robot for Vault, it may help you plan your configuration so that you use resources efficiently.

Robot Properties

| Property | Description |
|--------------|--|
| ACSLS Server | The name of the ACSLS server. StorageTek only. |
| ACS Number | The corresponding ACS number for this robot. You can obtain this information by using the Media Manager <code>tpconfig</code> or by using the ACSLS console commands such as <code>query acs all</code> or <code>query lsm all</code> . StorageTek only. |
| LSM Number | The corresponding LSM number for this robot. You can obtain this information by using the Media Manager <code>tpconfig</code> command or by using the ACSLS console commands such as <code>query acs all</code> or <code>query lsm all</code> . StorageTek only. |
| MAP Capacity | The capacity of the media access port (also known as cartridge access port). On StorageTek systems, you can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console. |
| MAP Numbers | The identifiers for the media access port. On StorageTek systems, you can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console. |

Methods of Configuration

You can use the NetBackup Administration Console to configure Vault. Alternatively, you can use the Vault Administration menu user interface on UNIX systems (invoked by the `vltaadm` command from a terminal window). These instructions describe using the NetBackup Administration Console to configure Vault.

In some circumstances, you may have to use Vault Administration menu interface to configure Vault, as follows:

- ◆ The NetBackup master server is installed on a UNIX host that does not support the NetBackup Administration Console and your computing environment does not have a system from which you can run the NetBackup Administration Console. (The NetBackup Administration Console that runs on UNIX systems is a Java-based application.)



- ◆ You have to connect to the UNIX system on which the NetBackup master server is installed from a remote system that does not have the NetBackup Administration Console. For example, if you have to connect to your network by using a dial-up connection over a telephone line, you may have to use a terminal window and use the Vault Administration interface.

For information about using the Vault Administration interface, see [“Using the Vault Administration Interface”](#) on page 189.

Note When you create or modify Vault configuration information, run only one instance of the NetBackup Administration Console or the Vault Administration interface. Using multiple instances at the same time may cause configuration information to be overwritten.

Configuring Vault Properties

Vault properties specify e-mail addresses for event notification and alternate media server names.

E-mail Tab

Use the Vault Properties **E-mail** tab to enter e-mail addresses for event notification and for reports:

- ◆ **E-mail address for notification of session status.** An e-mail notification is sent at the end of each vault session. The message provides a summary of the vault session, in the form of a `summary.log` file, and the status of the operation. The subject line of the e-mail message is formatted as follows:

Vault on *MasterServer*: Status *status_code* [*robot_number/vault_name/profile_name*]

By default, the e-mail is sent to the root or administrator user account on the system on which the NetBackup master server is installed. If you enter e-mail addresses in the **E-mail address for notification of session status** field, e-mail is sent to those addresses rather than to the root user. You cannot disable notification of session status.

- ◆ **Default e-mail address for reports.** You also can enter default e-mail addresses for report destinations. The e-mail addresses entered here are used in the E-mail address field on the Profile dialog **Reports** tab if you e-mail the reports (e-mailing reports is optional).

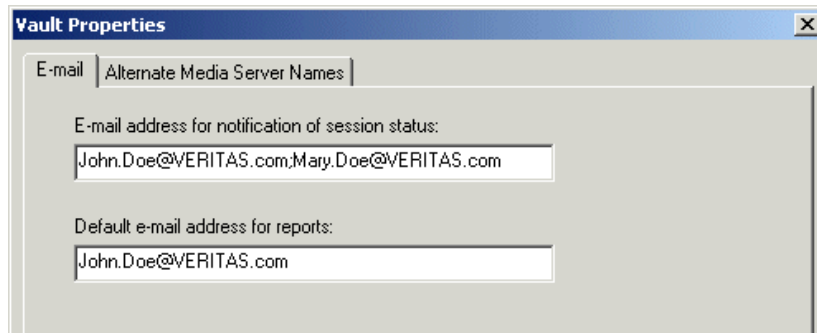
Related Topics

- ◆ [“Setting Up E-Mail”](#) on page 179



▼ To enter e-mail addresses

1. From within Vault Management, select **Vault Properties** from the **Actions** menu.
The Vault Properties dialog appears.
2. Select the **E-mail** tab.



3. Enter an e-mail address in the **E-mail address for notification of session status** field.
To enter more than one address, separate the addresses with commas.
4. Enter an e-mail address in the **Default e-mail address for reports** field.
To enter more than one address, separate the addresses with commas.
5. Click **OK**.

Alternate Media Server Names Tab

Use the Vault Properties **Alternate Media Server Names** tab to add alternate names of NetBackup media servers.

Adding alternate names for media servers simplifies configuration and helps ensure that all images eligible to be vaulted are chosen. Vault expands any occurrence of one of the names in a *server name group* to include all of the names in the group.

For every media server, you should add the fully qualified name, the short name, every name used by storage units that refer to it, any other names by which a media server has ever been known, and if you have multiple network interface cards (NICs) in the server, all server names or IP addresses associated with each NIC.



You also can create a server name group that includes different servers. Then, in the **Media Servers** field in the Profile **Choose Backups** tab, you only have to specify the server name group rather than the individual servers. This use of the Alternate Media Server Names dialog allows you to use one name to specify more than one server, which is useful if you want to duplicate images from multiple servers.

If you use the default, all media servers, for all of your vaults, you do not have to specify alternate media server names.

Alternative Media Server Names Background

A media server may have more than one name. For example, a server can have a fully qualified name, a short name, and more than one network interface card, each of which has its own name. If a media server has more than one storage unit, each storage unit can use a different name for that media server.

If a media server has more than one name, images backed up by it may be identified by an alternate name. If you specify only one of the names of that media server, images identified by the other names will not be vaulted.

Choose backups onfiguration is simplified if you specify media servers (that is, specify something other than the default, all media servers). If you add alternate media server names, you only have to specify one of those names in the **Media Servers** field of the Profile dialog **Choose Backups** tab; if you do not add alternate media server names, you must specify all the names associated with each media server on the **Choose Backups** tab.

Alternate Media Server Names Considerations

Be aware of the following associated with alternate media server names:

- ◆ You must have enough drives in the specified destination storage unit to keep up with the demand for duplication. If you do not, you risk a deadlock situation.
- ◆ The specified media servers must have access to the destination storage unit. If not, you risk a deadlock situation and your Vault job will fail. To prevent this situation, use the **Media Servers** criterion on the **Choose Backups** tab to ensure that only backups from certain media servers will be selected.
- ◆ If multiple duplication rules use different media server names that are part of a server name group, Vault processes only the first duplication rule; successive rules do not get processed. Also, because the media server name for the duplication rule is expanded to include all media server names in the group, all images written by all storage units that use those media server names are processed by the first duplication rule that uses any name from the group. All images are processed, but by the first duplication rule only.

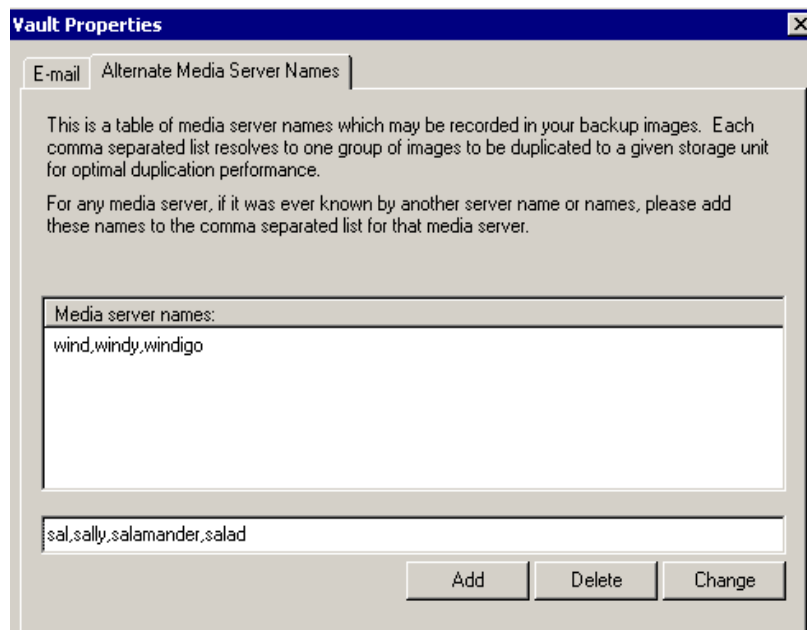
- ◆ Your configuration can send data over the network, depending on the media server(s) in use.

Note VERITAS recommends that you specify only one destination storage unit per server. If you specify more than one, you may create a problem because Vault does not have a mechanism to choose to which destination storage unit to send the duplicate images.

How to Add Alternate Media Server Names

▼ To add alternate media server names

1. From within Vault Management, select **Vault Properties** from the **Actions** menu.
By default, the **E-mail** tab of the Vault Properties appears.
2. Select the **Alternate Media Server Names** tab.



3. In the field below the Media Server Names window, enter all the alternate names for the media server, separated by commas, and then click **Add**.
 - ◆ To remove a media server name group you previously added, highlight it and click **Delete**.



- ◆ To change a name group you previously added, highlight it and click **Change**.
Each server name group should occupy one line in the Media Server Names window.

4. When finished, click **OK**.

Configuring Robots for Vault

Use the Vault Robot dialog to configure the robots from which Vault will eject media.

Vault Robot Dialog

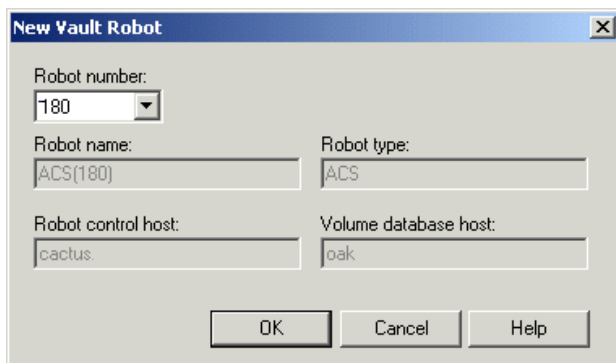
Use this dialog to specify robots that you want to use for your vault. Vault robots contain the media that has images to be stored off-site; that media will be ejected so it can be transferred to the vault. The images can be original images created during a backup job or duplicate images created by a Vault duplication job.

You can select any robots that are recognized by NetBackup and that have storage units associated with them. NetBackup assigns a number to each robot that it recognizes, and eligible robots are recognized by Vault.

▼ To configure a robot in Vault

1. In the NetBackup Administration Console, highlight **Vault Management**.
2. Open the **Actions** menu and select **New > New Vault Robot**.

The New Vault Robot dialog appears.



3. Select a robot number.

NetBackup Media Manager assigns a number to each robot that it recognizes, and eligible robots are recognized by Vault. Based on the robot number that you select, the other fields will be filled in automatically.

4. Click OK.

Robot Configuration Information

| Property | Description |
|----------------------|---|
| Robot Control Host | The name of the host that controls the robot. The robot control host is configured in Media Manager, and Vault uses that information to populate the Robot Control Host field. |
| Robot Name | The name the robot. The name is configured in Media Manager, and Vault uses that information to populate the Robot Name field. |
| Robot Number | The robot number assigned by Media Manager. NetBackup assigns a number to each robot that it recognizes, and eligible robots are recognized by Vault. |
| Robot Type | The robot type as configured in Media Manager. Vault uses that information to populate the Robot Type field. |
| Volume Database Host | The name of the Media Manager host that contains the Media Manager volume configuration information. This field is populated by Vault when you select a robot number. |

Creating a Vault

After you configure robots, you can create and configure vaults. Use the Vault dialog to configure vaults.

Vault Dialog

A vault is a logical entity that refers to a collection of removable media drives (usually tape drives) within a robot. You can use vaults to organize the data that is going off-site; for example, you can use one vault for payroll data and another vault for customer data.

If you are configuring a vault in an ACS robot, you also can configure the media access ports (MAPs) to use for eject operations.



The following is an example of the Vault dialog.

New Vault

Vault name:

Vault vendor:

Customer ID:

When vaulting, use:

☒ Slots for individual media

First off-site slot ID:

☐ Containers of many media

Geographies:

Robotic volume group:

Off-site volume group:

Media access ports to use:

Related Topics

- ◆ [“Volume Groups”](#) on page 45
- ◆ [“ACS MAP Overview”](#) on page 89

Best Practices

- ◆ [“Vaulting Paradigm”](#) on page 20
- ◆ [“Preferred Vaulting Strategies”](#) on page 20

Requirements for Creating a Vault

The following are the requirements for creating a vault:

- ◆ Robots must be configured in Vault.
- ◆ Each robot used for vaulting media requires at least one logical vault.
- ◆ A robot may contain multiple vaults, but a vault cannot span robots. Therefore, if you configured three TLD robots for Vault (not connected with pass-through devices), you must define at least three logical vaults, one for each TLD robot.
- ◆ Volumes in a vault must have the same density. If a robot has volumes of different density and you want to use all of those volumes for Vault, that robot must have a separate vault for each volume density.

How to Create a Vault

▼ To create a vault

1. In the NetBackup Administration Console, expand **Vault Management**.
2. Highlight a robot in the **Vault Management** tree.
3. From the **Actions** menu, choose **New > New Vault**.
The New Vault dialog appears.
4. Enter or select values for each field.
5. Click **OK**.

Vault Dialog Configuration Options

The following are the options you can configure in the Vault dialog:

Vault Dialog Configuration Options

| Property | Description |
|--------------------------|---|
| Change | For ACS robots only, the button used to configure media access ports for eject operations. If you click Change , the Media Access Ports dialog appears, in which you can add or remove MAPs from the Media Access Ports to Use list. |
| Containers of Many Media | Select if your media is stored in containers at your off-site storage location. |
| Customer ID | Your customer identification if you selected Iron Mountain as your vault vendor. You may have a separate customer ID for each logical vault. |



Vault Dialog Configuration Options (continued)

| Property | Description |
|----------------------------|--|
| First Off-site Slot ID | <p>The ID of the first slot in the off-site vault. This usually is provided by your vault vendor. Off-site slot IDs are often used by the vault vendor to track media. If your vendor does not use these identifiers, you can use the default first off-site slot ID of 1. Off-site slot IDs are unique only within a given vault.</p> <p>Slot IDs are assigned contiguously from the starting slot number. Ensure that the number of media in the vault does not exceed the range of slot IDs assigned by the vault vendor. With every session, Vault starts with the off-site slot ID and counts upwards, looking for slots that are no longer in use. Vault always fills in the gaps with newly vaulted media.</p> <p>In case multiple vaults are defined for the same vault vendor, you must divide the range of assigned slots between the various vaults. For example, if the vault vendor has assigned the range 1-2000 and you have defined 3 vaults for this vault vendor, then you can assign range 1-499 to vault 1, 500-999 to vault 2, and 1000-2000 to vault 3, assuming vault 3 has the maximum number of tapes to vault.</p> |
| Media Access Ports to Use | <p>For ACS robots only, the media access ports (MAPs) to use for media ejection for the current vault. To select or change MAPs to use, click Change. A Media Access Ports dialog appears, in which you can select the MAPs to use.</p> |
| Off-site Volume Group | <p>The name of the off-site volume group. The Off-site Volume Group indicates that media are in off-site storage. The name should describe the data, the vault vendor, the vault location, or a combination thereof so you can easily identify the volume group. Vault moves each piece of ejected media from the Robotic Volume Group into a standalone volume group (that is, a volume group that is not under the control of the robot). If the Off-site Volume Group does not exist, it will be created during the vault session. The off-site volume group name may contain up to 25 characters.</p> |
| Robotic Volume Group | <p>The name of the volume group associated with the robot for this vault. The Robotic Volume Group is the group that indicates media resides in a robot. Usually, NetBackup creates a robotic volume group when media are added to a robot.</p> |
| Slots for Individual Media | <p>Select if your media is stored in slots at your off-site storage location. If you select slots, you must complete the First Off-site Slot ID field.</p> |



Vault Dialog Configuration Options (continued)

| Property | Description |
|--------------|--|
| Vault Name | The name of the vault. The name should reflect its purpose. For example, if you are creating a vault primarily to duplicate and vault records from the finance department, you might call the vault Finance. The vault name may contain up to 25 characters. |
| Vault Vendor | The name of your off-site vault vendor (for example, Iron Mountain). If you select Iron Mountain, you also can configure Vault to put media lists into a file formatted in compliance with Iron Mountain's electronic processing specification. You can then send this file to Iron Mountain for electronic processing of the media lists. For more information about configuring Vault for Iron Mountain electronic processing, see "Configuring Reports" on page 91. |

Creating a Profile

After you configure vaults, you can create and configure profiles. Use the Profile dialog to configure profiles.

Profile Dialog

A Vault profile is a template for a vault job; it is a logical construct that contains the rules for selecting, duplicating, and ejecting media. A profile is associated with a specific vault, and at least one profile must exist for every vault. A vault can contain multiple profiles, although two profiles within the same vault cannot run simultaneously. Two different profiles can run simultaneously if each profile is in a different vault and if each profile uses a different off-site volume pool.

All profiles select images (that is, Choose Backups). You must select at least one of the following steps when you create a new Vault profile:

- ◆ Duplication
- ◆ Catalog Backup
- ◆ Eject

The other steps are optional so you can separate the Vault tasks into separate jobs if desired, using different jobs to accomplish different tasks. For example, you can use one job to select and duplicate images daily, and another job to eject media and generate reports weekly.

You can select or deselect any of these steps at any time during the configuration process.



Related Topics


- ◆ [“Running Multiple Sessions Simultaneously”](#) on page 98

How to Create a Profile

▼ To create a profile

1. Highlight a vault in the NetBackup Administration Console. From the **Actions** menu, choose **New > New Profile**.

The New Profile dialog appears.



2. In the **Name** field, type a name for the profile. VERITAS recommends that you use descriptive names.
3. Select the steps you want this profile to perform.

You must select at least one step. However, you can change the selections when you configure the profile. Because you must always configure the choose backups step, it is not displayed on this dialog.
4. Click **OK**.

The New Profile: *profile name* dialog appears.

Configuring a Profile

After you create a profile, the New Profile: *profile name* dialog appears. The New Profile dialog includes the following five tabs:

- ◆ The **Choose Backups** tab is where you specify the criteria for selecting backup images.



- ◆ The **Duplication** tab is where you configure duplication of the selected backup images.
- ◆ The **Catalog Backup** tab is where you choose how to back up the NetBackup and Media Manager catalogs. For efficient disaster recovery, you should vault a new catalog backup each time you vault data.
- ◆ The **Eject** tab is where you choose in which off-site volume pools Vault should look for the media you want to eject.
- ◆ The **Reports** tab is where you choose which reports to generate.

A profile must select images (Choose Backups). The other steps are optional so you can separate the tasks into separate jobs if desired, using different jobs to accomplish different tasks. For example, you can use one profile to select and duplicate images daily, and another profile to eject media and generate reports weekly.

▼ To configure a profile

1. If the Profile dialog is not displayed, highlight a profile in the NetBackup Administration Console window and select the **Change** icon in the toolbar.
2. Select the tab for each step that you are configuring and complete the fields.
3. When finished, click **OK**.

Configuring Choose Backups

The first step in configuring a profile is to specify the criteria by which Vault chooses backups. Rules that you specify on the Profile dialog **Choose Backups** tab determine which backups are chosen.

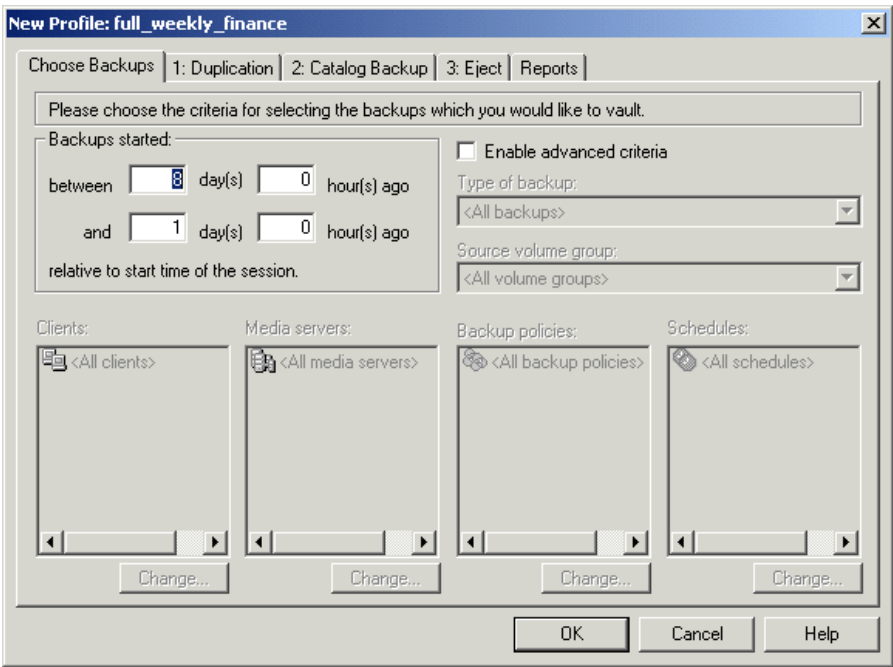
Choose Backups Tab

Use the **Choose Backups** tab to configure the search for images to be vaulted. The most basic criterion you can set is the time frame. To refine the search for images to vault, select **Enable Advanced Criteria** and then configure the advanced options. The most basic criterion you can set is the time frame; select the **Enable Advanced Criteria** box to specify advanced selection criteria.

Vault compares images in the NetBackup database with the criteria defined in the **Choose Backups** tab and generates a list of images that match the criteria. The image selection process will choose all images in the NetBackup catalog that match the criteria, even images that are in a different vault; the criteria that you specify on the other Profile dialog tabs determine whether Vault includes or excludes the images that are selected.



The following is an example of the Choose Backups tab:



Related Topics

- ◆ [“The List of Images to be Vaulted”](#) on page 104

Best Practices

- ◆ [“Overlap the Time Window in the Profile”](#) on page 22

Choose Backups Configuration Options

The following are the options you can configure in the **Choose Backups** tab:

Choose Backups Tab Configuration Options

| Property | Description |
|-----------------|--|
| Backup Policies | <p>A list of policies to use to select backup images. Enabled if you select Enable Advanced Criteria.</p> <p>To change the backup policies, click Change then choose the backup policies you want to include in the profile. Policies are based on the storage unit used for backups; because storage units are related to a specific robot number, choose the policies by robotic device.</p> |

Choose Backups Tab Configuration Options (continued)

| Property | Description |
|--------------------------|---|
| Backups Started | <p>The period of time from which the profile will select backups relative to the start time of the session.</p> <p>Time is expressed in terms of days and hours, relative to the time of the session. For example, assume the following settings:</p> <p style="padding-left: 40px;">between 8 day(s) 0 hour(s) ago and 1 day(s) 0 hour(s) ago</p> <p>If the session is started on October 12 at 1:00 pm, count backwards from October 12. The vaulted backups will be those started between October 4 at 1:00 pm (8 days before) and October 11 at 1:00 pm (1 day before).</p> <p>If you are selecting original backup images to send off site, the default time range is between 8 days and 1 day before the session runs; if you are duplicating images, the default time range is between 7 and 0 days.</p> |
| Change | Button used to display a dialog to change Clients, Media Servers, Backup Policies, or Schedules. |
| Clients | <p>The clients for which to select backup images. Enabled if you select Enabled Advanced Criteria.</p> <p>To change the clients, click Change then choose the clients you want to include in this profile.</p> |
| Enable Advanced Criteria | Option to enable advanced configuration. By default, All is selected for the advanced criteria options. |
| Media Servers | <p><i>Applies to NetBackup Enterprise Server only.</i></p> <p>The media servers from which to select backup images. Enabled if you select Enabled Advanced Criteria.</p> <p>To change the media servers, click Change then choose the media servers you want to include in this profile.</p> |
| Schedules | <p>A list of schedules to use to select backups. Enabled if you select Enabled Advanced Criteria.</p> <p>To change the default, click Change then choose the schedules you want to include in this profile. Schedules are based on the storage unit used for backups; because storage units are related to a specific robot number, choose the schedules by robotic device.</p> |



Choose Backups Tab Configuration Options (continued)

| Property | Description |
|---------------------|--|
| Source Volume Group | <p>A volume group from which to select backup images. Selecting a Source Volume Group restricts the search for images to those in that volume group rather than images in all volume groups. Usually, a Source Volume Group is specified if your master server has access to multiple robots and you want to duplicate images that reside on media in one robot to media in another robot. The images that are read are in the Source Volume Group in one robot; the images are written to media in the Robotic Volume Group in another robot.</p> <p>Volumes in the Source Volume Group will not be ejected. If you specify a Source Volume Group, you must configure duplication so that the source images are duplicated to media in the Robotic Volume Group, from which they will be ejected. <i>Exception:</i> if the Source Volume Group is the same as the Robotic Volume Group for the vault, volumes will be ejected.</p> <p>If you do not do duplication in Vault, you do not have to specify a Source Volume Group; if you specify a Source Volume Group, it has no effect on images that are vaulted.</p> |
| Type of Backups | <p>The types of backups (full, incremental, and so on) the profile will capture. Depending on the different types of backups you have configured in NetBackup policy management, you can choose the backup type. Only those types for which you have configured policies will be available for selection. If you want to vault all types of backups, accept the default. This is an optional criterion.</p> |

Configuring Duplication

Duplication reads original backup images that were created during a backup job and writes those images to other storage units. Rules you specify on the Profile dialog **Duplication** tab control image duplication. Duplication is optional; if you create multiple backup copies concurrently during a backup job and vault one of the originals, you do not need to duplicate images in Vault

Duplication Tab

Use the **Duplication** tab of the Profile dialog to configure the rules used to duplicate images. A duplication rule specifies the number of copies to create, a storage unit, off-site volume pool, retention period, media server (advanced configuration only), and what to do if an image copy fails (multiple copies only).

Duplication is optional; if you create multiple original backup copies concurrently during a backup job and vault one of the originals, you do not need to duplicate images in Vault.

Best Practices

- ◆ [“Avoid Resource Contention During Duplication”](#) on page 29
- ◆ [“Avoid Sending Duplicates Over The Network”](#) on page 35
- ◆ [“Increase Duplication Throughput”](#) on page 37
- ◆ [“Maximize Drive Utilization During Duplication”](#) on page 39

The Primary Backup Image

NetBackup restores from the primary backup image, and Vault duplicates from the primary backup image. (*Exception:* for performance reasons, Vault will duplicate from an image on disk rather than the primary backup if a disk image is available.) By default, the first original copy created during a backup job is the primary backup. Because both NetBackup and Vault use the primary backup, in most circumstances the primary copy should be the copy that remains in the robot.

If you choose to vault the original backup image, you can designate one of the duplicates that will remain in the robot as the primary backup.

If the primary backup image is off site, you cannot restore or duplicate the image until the media is injected into the robot or a local copy (if available) is promoted to primary. (*Exception:* Vault will duplicate from a disk image that is not the primary backup if a disk image is available.)

When the primary backup expires, NetBackup automatically promotes the backup copy that has the lowest number to primary.

Best Practices

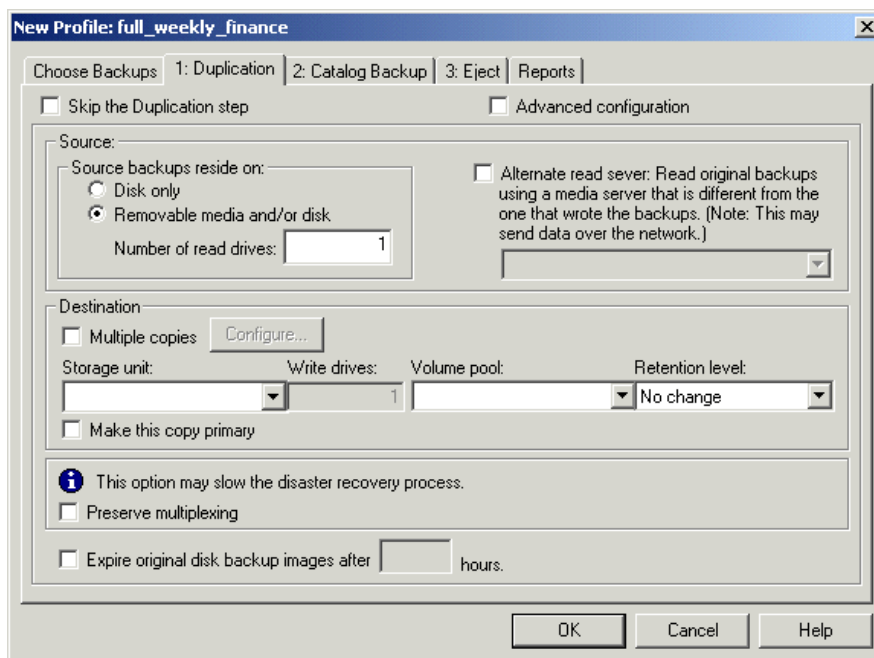
- ◆ [“Designate a Primary Copy and Keep It On Site”](#) on page 27

Basic Duplication

In basic duplication, you specify only one duplication rule. All backups are duplicated according to the same rule, and all selected images controlled by the specified master server are duplicated. You can create multiple copies of each backup image concurrently, but they are created using the same duplication rule.



The following shows the basic **Duplication** tab:



Advanced Duplication

Note Alternate read servers and multiple media servers apply to NetBackup Enterprise Server only.

Advanced duplication lets you specify more than one duplication rule. Vault determines which media server wrote each backup image and then applies the duplication rule corresponding to that media server to that image. In this context, the media server does not have any effect other than to identify which rule to apply to each image.

If a duplication rule does not specify an alternate read server, the media server that originally wrote the backup image will be used to read the original backup image during the duplication process.

Use advanced configuration only if you need to control exactly how to assign the backup images to be duplicated. The following may help you understand why to use advanced configuration:

- ◆ Your robot has different types of drives or media so that you have different storage units to use as destinations for the duplication process. In this case, you may want to balance the duplication job between multiple storage units. For example, you may

want to send the duplicate copies of all backup images written by one media server to a storage unit of one density and all backup images written by another media server to a storage unit of another density.

- ◆ Your profile is duplicating backup images to different media servers, each writing different types of data that require different retention periods. For example, if media server A backs up your customer database and media server B backs up warehouse inventory data, you may want to keep your customer database in off-site storage for a longer period of time (a different retention) than your warehouse inventory data.
- ◆ You have one media server that you need reserved for other operations. For example, you use multiple media servers for duplication but dedicate one media server for backups. For that one media server you would specify an alternate read server, and you would let the rest of the media servers handle their own duplication.

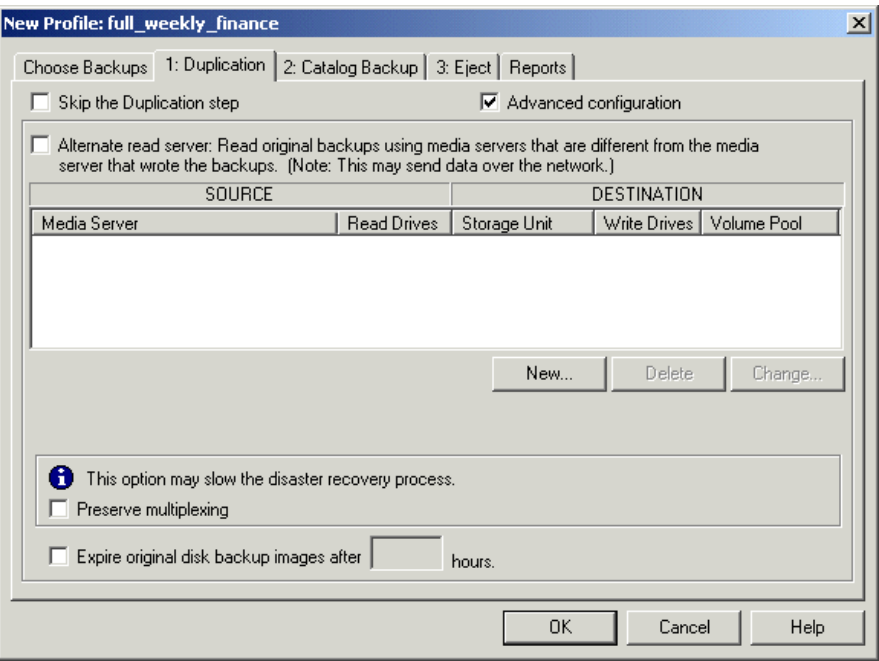
Note You do not need to configure advanced options if your profile duplicates images backed up by a single media server. To return to the basic **Duplication** tab, deselect the **Advanced Configuration** checkbox.

To avoid sending data over the network, do the following:

- ◆ For each duplication rule that does not specify an alternate read server, ensure that the media server controls both the source volumes and the destination storage units.
- ◆ For each duplication rule that specifies an alternate read server, ensure that:
 - ◆ The alternate read server is connected to all robots that have backup images written by the media server specified for this rule.
 - ◆ The alternate read server is the same server as the media server of the destination storage unit.



The following shows the **Duplication** tab when **Advanced Configuration** has been selected:



Duplication Tab Configuration Options

The following table describes configuration options for the **Duplication** tab.

Duplication Tab Configuration Options

| Property | Description |
|-----------------------|--|
| Alternate Read Server | <p><i>Applies to NetBackup Enterprise Server only.</i></p> <p>The name of an alternate read server. If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups. Using an alternate read server may transfer data over your network, affecting your site's computing environment. The Source Media Server and Alternate Read Server may be the same.</p> <p>By default this option is disabled. To configure an alternate read server, select Alternate Read Server then select a media server from the drop-down menu (or for advanced duplication, click New to configure duplication rules).</p> |

Duplication Tab Configuration Options (continued)

| Property | Description |
|------------------------------------|---|
| Change | <p>For advanced configuration only, the button used to display the Duplication Rule Dialog so you can change a destination media server and duplication rules for that server.</p> <p>If you selected Alternate Read Server on the Duplication tab, the Duplication Rule dialog will have fields for both Source Media Server and Alternate Read Server. If you did not select Alternate Read Server, only a Source Backup Server field appears.</p> |
| Configure | <p>For basic duplication only, the button used to display the Multiple Copies Dialog.</p> |
| Delete | <p>For advanced configuration only, the button used to delete the selected destination media server and duplication rules for that server.</p> |
| Expire Original Disk Backup Images | <p>The length of time (in hours) after the Vault session runs to expire the original backup on disk (applies only if original backup images are on disk). You can use this option to force an earlier expiration time for the images so the disk space is freed up for use by subsequent backups.</p> <p>If the duplication of a disk image is not successful, the disk image will not be expired.</p> |
| Make This Copy Primary | <p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image created during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.</p> |
| Multiple Copies | <p>Whether to create multiple copies concurrently. You can select Multiple Copies if the master server properties allow it. If you select Multiple Copies, click Configure to display the Multiple Copies Dialog. If you configure multiple copies, you cannot configure a Storage Unit, Volume Pool, Retention Level, or Primary Copy on the basic Duplication tab.</p> |



Duplication Tab Configuration Options (continued)

| Property | Description |
|-----------------------|--|
| New | <p>For advanced configuration only, the button used to display the Duplication Rule dialog, in which you can add a destination media server and duplication rules for that server.</p> <p>If you selected Alternate Read Server on the Duplication tab, the Duplication Rule dialog will have fields for both Source Media Server and Alternate Read Server. If you did not select Alternate Read Server, only a Source Backup Server field appears.</p> |
| Number of Read Drives | <p>The number of drives to use for reading backup images. When you enter a number of read drives, the same number will be entered into the Destination Write Drives field. You must have an equivalent number of read and write drives available.</p> |
| Preserve Multiplexing | <p>Whether to preserve multiplexing. Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process speeds up duplication, but slows down restores and disaster recovery processes. If the option to preserve multiplexing is selected, the multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session.</p> <p>If the source image is multiplexed and the Preserve Multiplexing option is selected, ensure that the destination storage unit configured for each copy has multiplexing enabled. Multiplexing is configured in NetBackup Management > Storage Units.</p> <p>Multiplexing does not apply to disk storage units or disk staging storage units as destinations. However, if the source is a multiplexed tape and the destination is a disk storage unit or disk staging storage unit, selecting Preserve Multiplexing ensures that the tape is read on one pass rather than multiple passes.</p> |

Duplication Tab Configuration Options (continued)

| Property | Description |
|---------------------------|---|
| Retention Level | <p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date will be the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1 (for more information, see “Assigning Multiple Retentions with One Profile” on page 121).</p> <p>When the retention period expires, information about the expired backup will be deleted from the NetBackup and Media Manager catalog, the volume will be recalled from off-site storage, and the backup image will be unavailable for a restore.</p> |
| Skip the Duplication Step | Select if you do not want to configure duplication. |
| Source Backups Reside On | The location of the backup images: disk or removable media or both. Vault will duplicate images from the primary backup images on removable media or from backup images on disk. |
| Storage Unit | <p>The name of a storage unit that contains the resources to which the copies of the backup images will be written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source and destination storage units can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, VERITAS suggests that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p> |



Duplication Tab Configuration Options (continued)

| Property | Description |
|--------------|--|
| Volume Pool | The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool will be ejected for transfer off-site. Do <i>not</i> use the volume pool that was used for the original backup; NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool. |
| Write Drives | The number of write drives. This value is the same as the number of read drives. |

Multiple Copies Dialog

The Multiple Copies dialog appears only if you select the **Multiple Copies** checkbox on the basic **Duplication** tab and then click **Configure**.

Use this dialog to create multiple copies of a backup image concurrently.

Multiple Copies

Copies:
4

All storage units must be connected to the same media server.

| | Primary: | Storage unit: | Write drives: | Volume pool: | Retention: | For each image if this copy fails: |
|---------|--------------------------|---------------|---------------|--------------|---------------------|------------------------------------|
| Copy 1: | <input type="checkbox"/> | <div></div> | <div>1</div> | <div></div> | <div>No chang</div> | <div>fail all copies</div> |
| Copy 2: | <input type="checkbox"/> | <div></div> | <div>1</div> | <div></div> | <div>No chang</div> | <div>fail all copies</div> |
| Copy 3: | <input type="checkbox"/> | <div></div> | <div>1</div> | <div></div> | <div>No chang</div> | <div>fail all copies</div> |
| Copy 4: | <input type="checkbox"/> | <div></div> | <div>1</div> | <div></div> | <div>No chang</div> | <div>fail all copies</div> |

OK

Cancel

Help

The following table describes configuration options for the Multiple Copies dialog.

Multiple Copies Dialog Configuration Options

| Property | Description |
|------------------------------------|---|
| Copies | <p>The number of copies to create concurrently. The number of copies you can choose cannot exceed the number of copies specified in the Maximum Backup Copies field for the NetBackup master server. (Configured in NetBackup Management > Host Properties > Master Server > <i>server_name</i> > Global NetBackup Attributes.) By default, the value is two: one original backup and one copy.</p> |
| For Each Image, If This Copy Fails | <p>The action to perform if a copy fails: Continue or Fail All Copies. In Vault, if you choose Fail All Copies, all copies <i>of that image</i> will fail, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault will again try to duplicate the image if the following conditions are true:</p> <ul style="list-style-type: none"> • The image is selected. • The Vault profile did not eject the primary backup. <p>By default, the option is configured to Fail All Copies in Vault. If you choose Continue for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted; it is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.</p> <p>For more information, see “Continue or Fail for Concurrent Copies” on page 144</p> |
| Primary | <p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image creating during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.</p> |



Multiple Copies Dialog Configuration Options (continued)

| Property | Description |
|--------------|---|
| Retention | <p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date will be the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1 (for more information, see “Assigning Multiple Retentions with One Profile” on page 121).</p> <p>When the retention period expires, information about the expired backup will be deleted from the NetBackup and Media Manager catalogs, the volume will be recalled from off-site storage, and the backup image will be unavailable for a restore.</p> |
| Storage Unit | <p>The name of a storage unit that contains the resources to which the copies of the backup images will be written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source and destination storage units can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, VERITAS suggests that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p> |
| Volume Pool | <p>The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool will be ejected for transfer off-site. Do <i>not</i> use the volume pool that was used for the original backup; NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool.</p> |
| Write Drives | <p>The number of write drives. This value is the same as the number of read drives.</p> |



Duplication Rule Dialog

The Duplication Rule dialog appears if you select **New** or **Change** on the **Advanced Configuration** options of the **Duplication** tab. If you selected **Alternate Read Server** on the **Duplication** tab, an **Alternate Read Server** option appears on the dialog.

Use the Duplication Rule dialog to create multiple copies of an image and to select different media servers and read servers for the copies.

The following table describes configuration options for the Duplication Rule dialog.

Duplication Rule Dialog Configuration Options

| Property | Description |
|-----------------------|--|
| Alternate Read Server | <p><i>Applies to NetBackup Enterprise Server only.</i></p> <p>The name of an alternate read server.</p> <p>If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups. Using an alternate read server may transfer data over your network, affecting your site's computing environment. The Media Server and Alternate Read Server may be the same.</p> <p>To configure an alternate read server, select a media server from the drop-down menu.</p> |



Duplication Rule Dialog Configuration Options (continued)

| Property | Description |
|------------------------------------|---|
| Copies | <p>The number of copies to create concurrently. The number of copies you can choose cannot exceed the number of copies specified in the Maximum Backup Copies field for the NetBackup master server. (Configured in NetBackup Management > Host Properties > Master Server > <i>server_name</i> > Global NetBackup Attributes.) By default, the value is two: one original backup and one copy.</p> |
| For Each Image, If This Copy Fails | <p>The action to perform if a copy fails: Continue or Fail All Copies. In Vault, if you choose Fail All Copies, all copies <i>of that image</i> will fail, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault will again try to duplicate the image if the following conditions are true:</p> <ul style="list-style-type: none">♦ The image is selected.♦ The Vault profile did not eject the primary backup. <p>By default, the option is configured to Fail All Copies in Vault.</p> <p>If you choose Continue for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted; it is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.</p> <p>For more information, see “Continue or Fail for Concurrent Copies” on page 144</p> |
| Media Server | <p><i>Applies to NetBackup Enterprise Server only.</i></p> <p>The name of the media server on which the backup images reside. The Media Server and Alternate Read Server may be the same.</p> |
| Number of Read Drives | <p>The number of drives to use for reading backup images. When you enter a number of read drives, the same number will be entered into the Destination Write Drives field. You must have an equivalent number of read and write drives available.</p> |

Duplication Rule Dialog Configuration Options (continued)

| Property | Description |
|--------------------------|--|
| Primary | <p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image creating during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.</p> |
| Retention | <p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date will be the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1 (for more information, see “Assigning Multiple Retentions with One Profile” on page 121).</p> <p>When the retention period expires, information about the expired backup will be deleted from the NetBackup and Media Manager catalog, the volume will be recalled from off-site storage, and the backup image will be unavailable for a restore.</p> |
| Source Backups Reside On | <p>The location of the backup images: disk or removable media or both. Vault will duplicate images from the primary backup images on removable media or from backup images on disk.</p> |



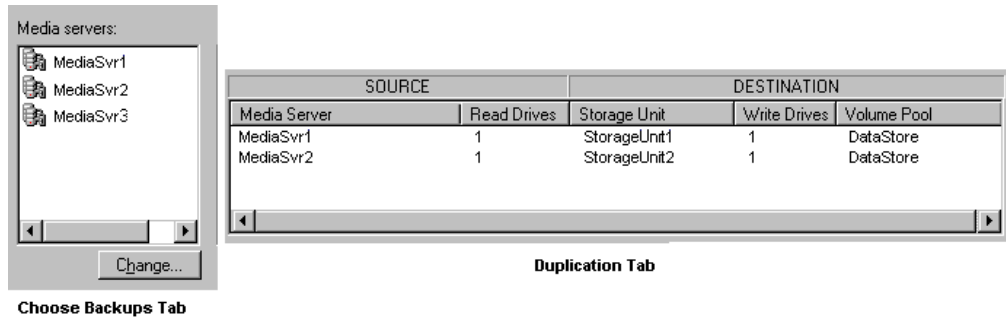
Duplication Rule Dialog Configuration Options (continued)

| Property | Description |
|--------------|---|
| Storage Unit | <p>The name of a storage unit that contains the resources to which the copies of the backup images will be written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source and destination storage units can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, VERITAS suggests that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p> |
| Volume Pool | <p>The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool will be ejected for transfer off-site. Do <i>not</i> use the volume pool that was used for the original backup; NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool.</p> |
| Write Drives | <p>The number of write drives. This value is the same as the number of read drives.</p> |

Treatment of Images Without Corresponding Duplication Rule

Note More than one media server applies to NetBackup Enterprise Server only.

In some cases, the profile may list more media servers in the Media Servers list on the **Choose Backups** tab (left) than in the advanced configuration view on the **Duplication** tab (right).



If this happens, images written by media servers that have no corresponding duplication rule must also be duplicated. Vault will duplicate those images but will try to minimize total duplication time by keeping as many drives as possible busy writing data until all images are duplicated. This is handled as follows:

- ◆ All images written by media servers that have a duplication rule are assigned to the appropriate duplication rule.
- ◆ As soon as one duplication rule has finished processing the images assigned to it, Vault will begin to assign images written by other media servers (media servers that have no rule of their own) to the duplication rule that had finished processing.
- ◆ As other rules complete the duplication of their assigned images, they too will be assigned images written by other media servers that have no rule of their own.
- ◆ Eventually all images written by all media servers listed on the **Choose Backups** tab will be duplicated and the duplication step will be complete. If you have more media servers listed on the **Choose Backups** tab than on the **Duplication** tab, there is only one way to ensure that large amounts of duplication data do not get sent over the network:
 - ◆ Every duplication rule must specify an alternate read server. For each duplication rule, the alternate read server must be the same as the media server of the destination storage unit(s).
 - ◆ All alternate read servers must be connected to all robots that have images written by any media server listed on the **Choose Backups** tab but not on the **Duplication** tab.

The previous configurations are best suited for a SAN environment where all media servers are visible to all robots.



Configuring Catalog Backup

The NetBackup and Media Manager catalogs consist of databases that contain information about the NetBackup configuration and any backups that have occurred, including records of the files backed up, the media on which the files are stored, and the media and storage devices that are controlled by Media Manager.


Catalog Backup Tab

Use the Profile dialog **Catalog Backup** tab to configure catalog backups. Vault creates a new catalog backup, it does not duplicate an existing NetBackup catalog backup.

The profile that creates the catalog backup must also eject the catalog backup media; because catalog backups are processed differently than data backups, Vault will only eject catalog media for the currently running profile. Because vaulting catalog media is so important, the catalog backup media will be ejected even if you do not specify the appropriate catalog volume pool in the Off-site Volume Pools list of the **Eject** tab (however, the eject step must be configured for the profile).

If no media are available in the dedicated off-site volume pool for catalog vaulting, the catalog backup will fail. If the catalog backup step fails but the remainder of the Vault job succeeds, the session will end with a partially successful (1) status. Data will be vaulted with no associated catalog backup. Vault does not retry the catalog backup because it can be a lengthy process that may exceed the time scheduled for the vault job (that is, successful data vault job plus unsuccessful catalog backup plus retried catalog backup probably will exceed the time window). VERITAS believes that it is better to vault the data without a catalog backup than to fail the job and potentially vault nothing at all for that session.

The following is the Catalog Backup tab:



Best Practices

- ◆ “[Vault NetBackup Catalogs](#)” on page 26

Default Catalog Locations

The following directories, in which the catalog database files reside, are selected by default for vaulting and listed in the Files to be Backed Up window of the **Catalog Backups** tab:

UNIX: `/usr/openv/netbackup/NetBackup/db`
 Windows: `install_path\NetBackup\db`

NetBackup scheduling information, error logs, and information about files backed up from client workstations.

UNIX: `/usr/openv/netbackup/volmgr/database`
 Windows: `install_path\volmgr\database`

The volumes, robots, and devices used in the current NetBackup configuration.

UNIX: `/usr/openv/netbackup/var`
 Windows: `install_path\var`

Information about licenses and authorization.



Catalog Backup Configuration Options

The following are the configuration options for the **Catalog Backup** tab:

Catalog Tab Configuration Options

| Property | Description |
|---|---|
| Backup Policy for Multiple-Tape Catalog Backups | <p>Large catalogs require more than one tape for a backup and also require a separate policy to support multiple-tape catalog vault operations. If your catalog backup requires more than one tape, you must specify the backup policy to use for the Vault catalog backup process.</p> <p>You can copy the NetBackup policy you use for multiple-tape catalog backup and modify it for use with Vault by specifying the dedicated multiple-tape catalog vault policy. For information about configuring NetBackup for large catalog backups, see “Protecting Large NetBackup Catalogs” in the <i>NetBackup System Administrator’s Guide, Volume I</i>.</p> |
| Change | The button used to change an entry in the Files to be Backed Up window. Highlight a catalog file and then click Change . You can change the pathname for the catalog, or you can click on the folder icon to browse and select a catalog. |
| Delete | The button used to delete an entry from the Files to be Backed Up window. Highlight a catalog file and then click Delete . |
| Files to be Backed Up | <p>The NetBackup database files to include in the catalog backup operation. By default, all the NetBackup and Media Manager catalog database files are included.</p> <p>If you select Use Catalog File Locations from NetBackup Catalog Backup Configuration, you cannot add, change, or delete pathnames to catalog files in the Files to be Backed Up window.</p> |
| Media Server | <p><i>Applies to NetBackup Enterprise Server only.</i></p> <p>The name of the media server that manages the catalogs. Only media servers that have a storage unit on the current Vault robot are displayed.</p> |
| New | The button used to add an entry in the Files to be Backed Up window. When you click New, a new blank line is created in the Files to be Backed Up window; you can enter a pathname for the catalog or you can click on the folder icon to browse and select a catalog. |

Catalog Tab Configuration Options (continued)

| Property | Description |
|--|--|
| Retention Period | <p>The number of days before the catalog backup expires and the volumes are recalled from the off-site vault.</p> <p>Vault recalls and reuses catalog backup media after the Retention Period has passed. Vault unassigns any catalog backup media that appears on the Picking List for Vault or Distribution List for Robot and returns that media to the catalog volume pool so it is available to reuse as catalog backup media. (Catalog backup media that were allocated from a scratch pool are not returned to the scratch pool because NetBackup never writes regular backup images onto catalog media, even expired Vault catalog backup media.)</p> |
| Skip the Catalog Backup Step | Select if you do not want to backup and vault the NetBackup and Media Manager catalogs. |
| Use Catalog File Locations from NetBackup Catalog Backup Configuration | Select if you want to use the catalog paths configured in NetBackup. Vault will use the paths specified on the Files tab of the NetBackup Catalog Backup dialog, and you cannot add, change, or delete pathnames in the Files to be Backed Up window. |
| Volume Pool | <p>The name of the volume pool that is reserved for off-site NetBackup catalog backups. You must use a dedicated off-site volume pool for catalog vaulting, and that volume pool must have unassigned media available (or in the scratch pool if used).</p> <p>Because vaulting the catalog backup is a critical step, Vault ejects the catalog backup volume even if you do not specify this volume pool in the eject step.</p> |

Configuring Eject

Rules that you specify on the Profile dialog **Eject** tab determine how media are ejected.

Eject Tab

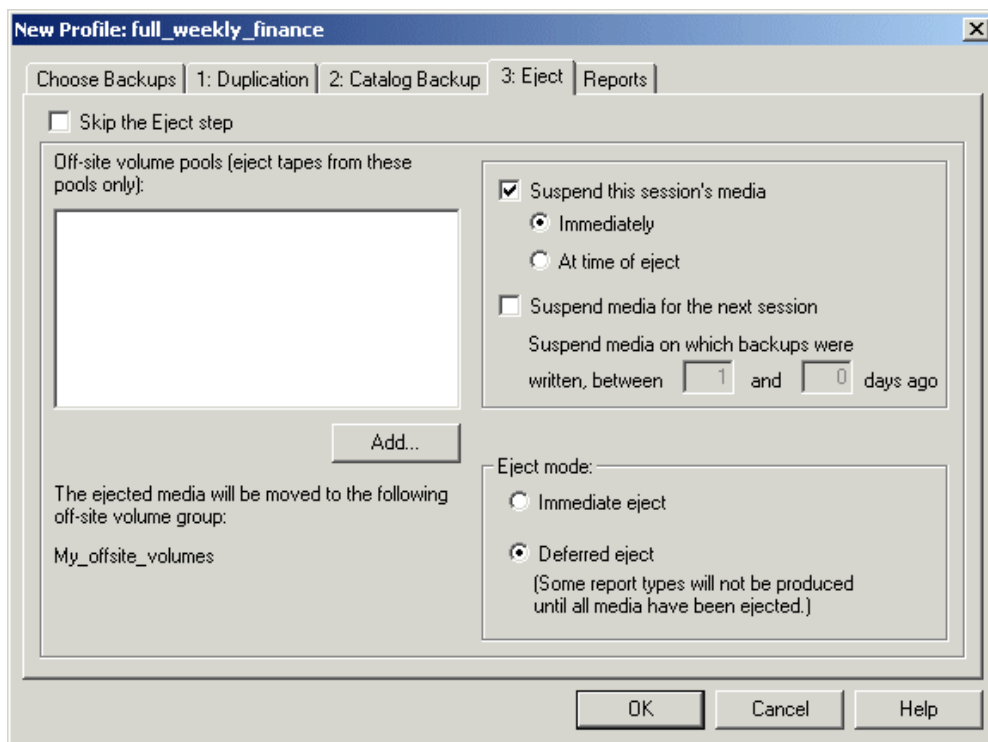
Use the Profile dialog **Eject** tab to specify the volume pools from which to eject media, the suspend options, and whether to eject the media immediately when the profile runs or later.

By default, Vault suspends media when it is ejected. To change the default suspend behavior, select different suspend options. Vault only suspends the volumes in volume pools specified in the off-site volume pools list.



If you create catalog backup media in a profile in which you eject media, that catalog media will be ejected automatically; you do not have to specify the appropriate catalog volume pool in the off-site volume pools list of the **Eject** tab. Catalog backups are processed differently than data backups; therefore, Vault only ejects catalog media for the current vault profile.

The following is the **Eject** tab of the Profile dialog.



Related Topics

- ◆ [“Ejecting Media”](#) on page 106

Best Practices

- ◆ [“Specify Robotic Volume Group When Configuring a Vault”](#) on page 23
- ◆ [“Avoid Vaulting Partial Images”](#) on page 24
- ◆ [“Suspend Vaulted Media”](#) on page 28
- ◆ [“Defer Ejection”](#) on page 29

Eject Configuration Options

The following are the configuration options for the **Eject** tab:

Eject Tab Configuration Options

| Property | Description |
|---|--|
| Add | The button used to add a volume pool to the eject list. If you click Add , the Volume Pools dialg appears, in which you can add or remove volume pools from the eject volume pool list. |
| At Time of Eject (Suspend this Session's Media) | Suspend the media when it is ejected. If you also select deferred eject, images can be written to the media until it is ejected. Select this option if you want the media sent off-site to be full. Suspend at time of eject is the default. |
| Deferred (Eject Mode) | Defer media ejection until a later time. The reports that are marked with an asterisk (*) on the Reports tab are generated only when all media selected by the profile have been ejected. |
| Immediate (Eject Mode) | Eject media immediately as part of the current Vault job. The reports that are marked with an asterisk (*) on the Reports tab are generated only when all media selected by the profile have been ejected. |
| Immediately (Suspend this Session's Media) | Suspend the media during the current session. No more images will be written to the media even if ejection is deferred. |
| Off-site Volume Pools | The names of the volume pools from which to eject media. Only the media in the pools that contain images that meet the selection criteria are ejected. (Catalog backup media for the profile are ejected even if you do not specify the catalog volume pool.) If you use a <code>vlt_ejectlist_notify</code> script to eject media not created by NetBackup or Vault, you must add the volume pool in which that media resides to the Off-site Volume Pools list of the profile that you run to eject that media. For information about notify scripts, see “Using Notify Scripts” on page 131. |
| Skip the Eject Step | Select if you do not want to eject media with this profile. |



Eject Tab Configuration Options (continued)

| Property | Description |
|--|---|
| Suspend Media for the Next Session (Suspend Media on Which Backups Were Written) | <p>Select to suspend original backup media, then enter the number of days before the Vault job to suspend media.</p> <p>Use this option only if you vault original images and want to prevent NetBackup from writing partial images on backup media.</p> <p>You should carefully consider whether to use this option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again to select media to suspend. Also, this option will not suspend media that is in use, such as media to which NetBackup is writing backup images. Therefore, some partial images on vaulted media may be acceptable.</p> <p>This option will suspend duplicate media created by Vault; however, the Suspend this Session's Media option is a better choice for suspending duplicate media because it does not use CPU cycles to select media to suspend.</p> <p>For more information about preventing partial images from being vaulted, see "Avoid Vaulting Partial Images" on page 24.</p> |
| Suspend this Session's Media | <p>Select to suspend media in the eject list, then select either Immediately or At Time of Eject.</p> <p>Suspend at time of eject is the default.</p> |

Media Ejection Overview

During media ejection operations, Vault moves the media to be ejected into the default media access port (MAP) of the robotic library. You must extend the MAP, remove the media, and then retract the MAP. If more media will be ejected, Vault will continue to fill the MAP until all media are ejected. For libraries that have separate MAP doors such as libraries connected via pass-through mechanisms, all doors are treated as one continuous MAP by NetBackup. In other words, each time you are prompted by NetBackup, open all the doors, empty all the MAPs, and then close all the doors.

For ACS robots that have multiple MAPs, you can specify the MAPs to eject media to when you configure the robot in Vault.

If you use a library that does not have a MAP, you must remove the media from the library slots; however, you also have to perform the eject operation in Vault so that the appropriate database entries are completed. To eject media from a library that does not have a MAP, do the following:

- ◆ Configure the profiles for deferred eject.
- ◆ Eject the media manually (see "[Ejecting Media](#)" on page 106).

- ◆ Remove the media from the library slots.

Do not inventory the robot until you remove the media from the MAP or library slots. If you do, you will have to revault the media.

ACS MAP Overview

Applies to NetBackup Enterprise Server only.

Automated cartridge system (ACS) robots can have multiple library storage modules (LSMs), each with multiple media access ports (MAPs). When you configure a vault that uses an ACS robot, you can specify any MAP or a subset of MAPs to be used for media ejection. Vault will eject media to as few of the configured MAPs as possible based on a nearest MAP algorithm. The algorithm considers the volumes to be ejected, the MAPs configured for ejecting in the vault, and the configuration of the LSMs. The algorithm assumes that the LSMs are connected in a line; if your LSMs are connected in a configuration other than a line, see “Adjacent LSM Specification for ACS Robots” and “Media Access Port Default for ACS Robots” in the *NetBackup Media Manager System Administrator’s Guide*.

Any MAP does not mean *all* MAPs; media will not be ejected to all MAPs, media will be ejected to the nearest MAP in each LSM. If you specify any MAP:

- ◆ MAPs that have only one element will not be used.
- ◆ Vault will choose from MAPs that are on-line when the eject begins; MAPs that are off-line are not considered for eject operations.
- ◆ If only a subset of MAPs are used during ejection, all MAPs will be busy and unavailable (for example, if the media are ejected to only two MAPs in one LSM, all MAPs will still be busy).

For all other robot types that have MAPs, media are ejected to the default MAP. NetBackup does not support ejecting to multiple MAPs for other robot types.

Eject Mode (Immediate or Deferred)

You can eject media immediately when the profile runs or defer ejection until later. If you use one profile to choose and duplicate images daily and another profile to eject the media, you should specify deferred eject for the profile that selects and duplicates images and immediate eject for the profile that ejects media. If you defer eject, you should also defer reports.

If you select deferred eject, other actions are required to eject the media for the session. The following are the methods you can use to eject media after the session has ended:

- ◆ Eject for one session only, as follows:
 - ◆ Use the NetBackup Administration Console to eject media for the session.



- ◆ Use the Vault Operator Menu to eject media for the session.
- ◆ Use the `vlteject` command to eject media for the session.
- ◆ Create a Vault policy and enter the appropriate `vlteject` command and options in the file list.
- ◆ Eject for multiple sessions for a specific profile, as follows:
 - ◆ Configure a Vault profile to duplicate only, and configure a Vault policy to run this profile on the days you want to select and duplicate images.
 - ◆ Configure a second Vault profile to do the catalog backup and eject steps. This profile should use the same image selection criteria as the profile that duplicates images. Configure a Vault policy to run this profile on the day you want the media ejected.

This method for duplicating and ejecting media provides the added benefit of consolidated reports that are not organized by session.

- ◆ Eject for all sessions for a specific vault or for all sessions for all vaults (consolidated eject) by doing one of the following:
 - ◆ Use the NetBackup Administration Console.
 - ◆ Use the Vault Operator Menu.
 - ◆ Use the `vlteject` command.
 - ◆ Create a Vault policy and enter the appropriate `vlteject` command and options in the file list.

If you defer eject operations, you should also defer reports until you eject media.

Related Topics

- ◆ [“Ejecting Media”](#) on page 106



Media Ejection Timeout Impact

The media ejection timeout period is the amount of time the eject process will wait for the removal of the ejected media from the media access port (MAP) before the media is injected back into the robot and an error condition occurs. The timeout period varies depending on the capability of each robot. The following table shows the ejection timeout periods for robots.

Media Ejection Timeout Period for Vault

| Robot | Timeout Period | Note |
|----------------------------------|----------------|---|
| Automated Cartridge System (ACS) | One week | Applies to NetBackup Enterprise Server only |
| Tape Library 8MM (TL8) | One week | |
| Tape Library DLT (TLD) | One week | |
| Tape Library Half-inch (TLH) | None | Applies to NetBackup Enterprise Server only |
| Tape Library Multimedia (TLM) | One week | Applies to NetBackup Enterprise Server only |
| Robots that do not have MAPs | None | |

For robots that do not have timeout periods or do not have MAPs, select deferred eject and then eject the media manually. When you eject the media, ensure that the media are removed from the library in a timely manner so that other operations can occur.

Status messages displayed by the NetBackup Administration Console or on the command line provide information about the status of inject, eject, or inventory operations.

Caution If media are not removed and the timeout period expires, the Vault reports will not accurately show the status of the media. To recover, you should use the Vault Operator Menu (`vltopmenu`) or `vlteject` to eject the media that was not removed from the library and generate the reports.

Configuring Reports

Options you select on the Profile dialog **Reports** tab determine which reports are generated, when they are generated, and how and to whom they are distributed.



You and your off-site storage vendor can use the reports to determine which media should be moved between your site and the off-site storage location and the timing of the moves.

Reports Tab

Use the **Reports** tab to select which reports to generate for the profile, where to distribute them, and when to generate them (immediately when the profile runs or deferred until later). You also can change the titles of the reports. Generating reports is optional.

Reports can be generated for one session or for multiple sessions (known as *consolidating* your reports and ejections).

Note You must specify a report destination in the profile **Reports** tab so that reports will be generated.

The following is the **Reports** tab of the Profile dialog.

New Profile: full_weekly_finance

Choose Backups | 1: Duplication | 2: Catalog Backup | 3: Eject | **Reports**

Report header:

Reports for media going off-site:

- ☒ Picking List for Robot
- ☐ Distribution List for Vault
- ☐ Detailed Distribution List for Vault
- ☐ Summary Distribution List for Vault

Reports for media coming on-site:

- ☒ Picking List for Vault (*)
- ☐ Distribution List for Robot (*)

Recovery Report for Vault (*)
between and days ago

Inventory reports:

- ☒ Vault Inventory (*)
- ☐ Off-site Inventory (*)
- ☐ All Media Inventory (*)

Exception reports:

- ☐ Non-vaulted Images

Report destination:

- ☒ E-mail:
- ☒ Print command:
- ☐ Folder:
- ☒ Iron Mountain FTP file
- Destination folder:

Report mode:

Reports marked (*) will not be produced until all media have been ejected.

- ☐ Immediate reports
- ☒ Deferred reports

Related Topics

- ◆ [“Generating Reports”](#) on page 159



- ◆ “[Consolidating Reports](#)” on page 162
- ◆ “[Vault Report Types](#)” on page 164
- ◆ “[Setting Up E-Mail](#)” on page 179

Best Practices

- ◆ “[Ensure Report Integrity](#)” on page 40
- ◆ “[Generate the Lost Media Report Regularly](#)” on page 41

Reports Configuration Options

The following are the **Reports** tab configuration options.

Reports Tab Configuration Options

| Property | Description |
|--------------------------------------|--|
| All Media Inventory | Select to generate the All Media Inventory report. |
| Change Report Titles | <p>To change the title of a report, click Change Report Titles to display the Change Report Titles dialog</p> <p>If you change a title, the new title will be displayed on the Reports tab rather than the original names.</p> <p>If you consolidate your reports and also change titles, use the same title for all profiles whose reports will be consolidated. The report title is printed on the reports and appears in the e-mail subject line if you e-mail the reports.</p> |
| Deferred Reports | <p>Defer generating the reports until after the session has completed (for example, if you run Vault sessions daily and eject media weekly). Deferred is the default.</p> <p>Reports marked with an asterisk (*) are generated only when all media selected by the profile are ejected.</p> |
| Detailed Distribution List for Vault | Select to generate the Detailed Distribution List for Vault report. |
| Directory (UNIX systems) | <p>Enter the pathname of a directory in which you want the reports saved. The file name of the report will include the session ID, so you can use the same folder for all Vault reports.</p> <p>By default, Vault reports are saved in each session directory. Reports will still be saved in the session directories even if you specify a directory for reports.</p> |



Reports Tab Configuration Options (continued)

| Property | Description |
|-----------------------------|---|
| Distribution List for Robot | Select to generate the Distribution List for Robot report. |
| Distribution List for Vault | Select to generate the Distribution List for Vault report. |
| E-mail | <p>Enter e-mail addresses, separated by commas, semicolons, or spaces. If you have already entered this information on the E-mail tab of the Vault Properties dialog, it will show up here automatically when you select the E-mail checkbox.</p> <p>If you e-mail reports, you should ensure that e-mail is set up correctly in NetBackup; see “Setting Up E-Mail” on page 179.</p> |
| Folder (Windows systems) | <p>Enter a pathname of a folder in which you want the reports saved. The file name of the report will include the session ID, so you can use the same folder for all Vault reports.</p> <p>By default, Vault reports are saved in each session directory. Reports will still be saved in the session directories even if you specify a directory for reports.</p> |
| Immediate Reports | Generate the reports immediately as part of the current vault session. Reports marked with an asterisk (*) are generated only when all media selected by the profile are ejected. |
| Iron Mountain FTP File | If you selected Iron Mountain as your vault vendor (in the New Vault dialog), Iron Mountain FTP file and Destination folder items are displayed. If you want a file generated that you can send by suing FTP to Iron Mountain, select Iron Mountain FTP file and enter the name or browse to choose the name of the Destination folder to which the file will be written. Sending the file to Iron Mountain is not part of the vault process. |
| Non-vaulted Images | Select to generate the Non-vaulted Images report. |
| Off-site Inventory | Select to generate the Off-site Inventory report. |
| Picking List for Robot | Select to generate the Picking List for Robot report. |
| Picking List for Vault | Select to generate the Picking List for Vault report. |



Reports Tab Configuration Options (continued)

| Property | Description |
|-------------------------------------|---|
| Print Command | <p>Select to print reports, then enter a print command for a printer on the master server on which Vault is installed. Specify the full path to the print command. If you want to add an alternate print command to print the Recovery Report in landscape format, separate the print commands with a plus (+).</p> <p>You should test the print commands from a command line on the server on which Vault is installed to ensure that they work correctly.</p> <p>On Windows systems, the print server must grant the appropriate access permissions to the individual users of the vault commands, and also to the SYSTEM user in the LocalSystem account on the machine from which the print request originates.</p> |
| Recovery Report for Vault | Select to generate the Recovery Report for Vault, then enter the time period for the report. |
| Report Header | If you want certain text to appear at the top of every report, enter it in the Report Header box. The header will appear on all reports. |
| Summary Distribution List for Vault | Select to generate the Summary Distribution List for Vault report. |
| Vault Inventory | Select to generate the Vault Inventory report. |

Report Mode (Immediate or Deferred)

Similar to the eject mode, you can specify whether reports should be generated immediately when the profile runs or deferred until later. If you defer eject, you should also defer reports. If you defer reports, you must perform or schedule another action to generate the reports.

Because some reports are generated only when media are ejected, you may choose to defer reports until the media are ejected. For example, if you duplicate images daily and eject media weekly, you can defer reports for the profile that duplicates images and use the profile that ejects media to generate reports.



If the corresponding eject process has been completed when you generate reports, all pending reports are generated and distributed; those reports will not be regenerated if you run deferred reports again. If eject has not been completed, the subset of reports that do not depend on completion of eject will be generated; those reports will be generated again if deferred reports are run again.

If you defer reports from multiple sessions and then generate them together, it is known as consolidating reports.

Reports that Depend on Eject

Reports marked with an asterisk (*) on the profile dialog **Reports** tab are generated only when all media selected by the profile are ejected:

- ◆ Reports for media coming on-site:
 - ◆ Picking List for Vault
 - ◆ Distribution List for Robot
- ◆ Inventory reports:
 - ◆ Vault Inventory
 - ◆ Off-site Inventory
 - ◆ All Media Inventory)
- ◆ Recovery Report for Vault

Vaulting and Managing Media

The following provide information about vaulting and managing media:

- ◆ [“Running a Vault Session”](#) on page 98
- ◆ [“Previewing a Vault Session”](#) on page 100
- ◆ [“Stopping a Vault Session”](#) on page 101
- ◆ [“Resuming a Vault Session”](#) on page 101
- ◆ [“Monitoring a Vault Session”](#) on page 102
- ◆ [“The List of Images to be Vaulted”](#) on page 104
- ◆ [“Ejecting Media”](#) on page 106
- ◆ [“Injecting Media”](#) on page 112
- ◆ [“Vaulting and Managing Media in Containers”](#) on page 115
- ◆ [“Assigning Multiple Retentions with One Profile”](#) on page 121
- ◆ [“Vaulting Additional Volumes”](#) on page 124
- ◆ [“Revaulting Unexpired Media”](#) on page 126
- ◆ [“Tracking Volumes Not Ejected by Vault”](#) on page 128
- ◆ [“Vaulting Media Not Created by NetBackup”](#) on page 129
- ◆ [“Notifying a Tape Operator When Eject Begins”](#) on page 130
- ◆ [“Using Notify Scripts”](#) on page 131
- ◆ [“Clearing the Media Description Field”](#) on page 134
- ◆ [“Ensuring Available Media for Catalog Backups”](#) on page 134
- ◆ [“Deassigning Vaulted NetBackup Catalog Media”](#) on page 135
- ◆ [“Restoring Data from Vaulted Media”](#) on page 136
- ◆ [“Replacing Damaged Media”](#) on page 137



Running a Vault Session

A Vault session, or vaulting job, is the process of executing the steps specified in a Vault profile. Before you can run a Vault session, at least one robot, one vault, and one profile must be configured. Usually, Vault sessions are run automatically by the NetBackup scheduler; the scheduler runs the commands in a Vault policy

For more information about running a Vault session, see the following:

- ◆ [Running Multiple Sessions Simultaneously](#)
- ◆ [Running a Session Automatically](#)
- ◆ [Running a Session Manually](#)

You can also run a vault session by using the Vault Administration menu interface (UNIX systems only). For more information, see “[Using the Vault Administration Interface](#)” on page 189.

Running Multiple Sessions Simultaneously

Two or more profiles within the same vault cannot run simultaneously. However, it is possible to run more than one Vault session at the same time with the following restrictions:

- ◆ Each profile must be in a different vault.
- ◆ Each profile must use a different on-site volume pool than any of the other profiles that will run at the same time.
- ◆ Each profile must use a different Off-site Volume Pool than any of the other profiles that will run at the same time.
- ◆ Each profile must use a different Off-site Volume Group than any of the other profiles that will run at the same time.
- ◆ Only one Vault catalog backup may be active at any time. Therefore, configure your profiles so that only one catalog backup at a time is performed by Vault.
- ◆ Your environment must have enough resources so that the sessions can run simultaneously. For example, the number of drives required for simultaneous read and write operations for all profiles that will run simultaneously should not exceed the number of drives available.

Other factors can affect whether more than one profile can run at the same time successfully, including the capability of robotic libraries to perform multiple operations at the same time. Therefore, you should consider other actions and practices that can improve the success of simultaneous sessions, such as the following:

- ◆ Defer the eject process for all simultaneous sessions.

- ◆ Reserve a specific time period for Vault media ejection.
- ◆ Eject media by configuring a profile for ejection only or ejecting media manually.
- ◆ Do not inject or eject other media while Vault is ejecting media.
- ◆ Do not inventory a robot while Vault is ejecting media.

For information about the capabilities of the robotic libraries supported by NetBackup, see the *NetBackup Media Manager System Administrator's Guide*.

Running a Session Automatically

To run a vault session automatically, configure a NetBackup policy that runs a Vault profile according to the schedule in the policy.

For instructions to create a Vault policy that runs a session automatically, see “[Vault Policies](#)” on page 46.

Running a Session Manually

You can run a Vault session manually either by using the NetBackup Administration Console to initiate the session or by invoking the `vlrun` command from a command line.

Running a Session from the Administration Console

To run a Vault session from the NetBackup Administration Console, you either invoke the policy manually or invoke the profile manually.

▼ To invoke a Vault policy

1. Expand **NetBackup Management > Policies** in the left pane of the Administration Console window.
2. Select the policy you want to run.
3. Select **Actions > Manual Backup**.

The policy will be run at that time and also at its scheduled time and date.

▼ To invoke a Vault profile

1. Expand **Vault Management** in the left pane of the Administration Console window.
The names of the robots configured for NetBackup appear.



2. In the left pane, expand the robot that contains the vault and profile you want to run.
3. In the left pane, select the vault that you want to run.
4. In the details (right) pan, click on the profile you want to run.
5. Select **Actions > Start Session**.

Start Session remains highlighted until the session begins. When the session starts, the Console displays a message similar to the following:

```
Manual vault session for profile has been started. Use the Activity  
Monitor to view progress.
```

By default, the Details Pane of the Administration Console window does not show the Volume Pools (Ejected) and Report Destination columns. You can add, delete, or rearrange the columns displayed in the Details Pane by selecting **View > Columns > Layout**.

Running a Session from a Command Line

To run a vault session from a command line, first add the path in which the NetBackup executable files are installed to your PATH environment variable. Then, invoke `vltrun` from a command line, specifying the robot number, vault, and profile as in the following example:

```
vltrun robot_number/vault_name/profile_name
```

Alternatively, you can specify only the profile if it has a unique name.

For information about the `vltrun` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Previewing a Vault Session

Before you run a Vault session, you can preview the session to verify that the profile selects the appropriate images for off-site storage. To preview a session, use the `vltrun` command with the `-preview` option, specifying the robot number, vault, and profile as in the following example:

```
vltrun robot_number/vault_name/profile_name -preview
```

Alternatively, you can specify only the profile if it has a unique name.

The `vltrun -preview` option starts a new vault job, performs a search on the image catalog based on the criteria specified on the Choose Backups tab, writes the names of the images to a `preview.list` file, and then exits. Vault does not act on the images selected.

After you run the preview option, check the results in the `preview.list` file, which is located in:

- ◆ UNIX: `/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx`
- ◆ Windows: `install_path\NetBackup\Vault\sessions\vault_name\sidxxx`

Under certain circumstances, the `preview.list` file may contain more backup images than will be vaulted:

- ◆ If the profile is configured to duplicate only disk images, selected images on removable media will not be vaulted.
- ◆ If images in the list do not have a copy on media in one of the Off-site Volume Pools listed for the eject step, they will not be vaulted.

Stopping a Vault Session

You can use Activity Monitor to stop a Vault session. The Activity Monitor must be configured to show the Vault fields.

▼ To stop a vault session

1. In the Activity Monitor, highlight the vault session you want to stop.
2. From the **Action** menu, select **Cancel Job**.

Note If a vault session fails, you cannot run a new session until the old session has ended. Use **Cancel Job** to end the failed session.

Resuming a Vault Session

If your vault job fails or returns a Partially Successful completion status, you should first consult the NetBackup Administration Console Activity Monitor or the notification of session status (the session's `summary.log`). If they do not provide enough information to determine the cause of the problem, examine the other log files (see “[Debug Logs](#)” on page 202).

After you determine the cause, you can do one of the following:

- ◆ If the session had reached the Eject step or had attempted to generate reports before encountering problems, you can use `vltopmenu` (or `vlteject`) to finish the eject and/or reporting for the session.



- ◆ Start a new session for your profile. If you are doing duplication, Vault will not duplicate any images it already duplicated, but it will eject those images if the profile is configured to eject. (This behavior is similar to checkpoint and restart.)

Monitoring a Vault Session

If you configure the NetBackup Administration Console Activity Monitor to display the Vault fields, you can use the Activity Monitor to monitor the progress of Vault jobs. For a Vault job initiated by the NetBackup scheduler, the **Policy** field displays the policy name. If the Vault job is invoked by any means other than the NetBackup scheduler, the **Policy** field is empty.

For information about configuring the Activity Monitor to display fields other than the default, see the “Monitoring NetBackup Activity” section in the *NetBackup System Administrator’s Guide, Volume I*. The following are the fields that display Vault job attributes in the Activity Monitor:

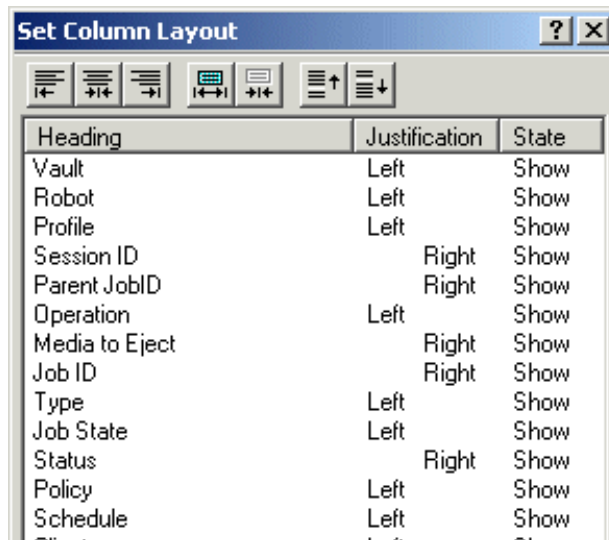
Vault Job Activity Monitor Fields

| Field | Description |
|----------------|--|
| Media to Eject | The number of tapes to be ejected for a vault session. Note If the Vault profile is configured for manual eject, the tapes may not have been ejected yet. |
| Operation | For vault jobs, the field contains one of the following values. These values progress from the first value to the last as the vault job progresses: Choosing Images Duplicating Images Choosing Media Catalog Backup Eject and Report |
| Parent JobID | A Vault job that duplicates images starts one or more <code>bpduplicate</code> processes. Each of these child jobs refers to the job ID of the Vault job (parent) that started it. |
| Profile | The name of the profile that holds the processing information for a vault session. |
| Robot | The name of the robot the vault is associated with. |

Vault Job Activity Monitor Fields (continued)

| Field | Description |
|------------|--|
| Session ID | The unique numeric value that identifies the vault session. Session ID assignment starts at 1 the first time a vault session is run after vault has been installed. The value is incremented by one every time a new vault session runs. |
| Vault | The name of the vault under which this session is running. |

The following is an example of the Activity Monitor column layout dialog showing the Vault fields at the top of the window:



Extended Error Codes

Vault jobs may exit with exit-status values greater than 255. These values are called extended error codes because they extend beyond the standard 255 NetBackup error codes. If a vault job exits with an extended error code, the exit status returned to the shell is 252. NetBackup has adopted the convention that the exit status 252 means that an extended error code is returned via stderr, in this message:

```
EXIT status = extended error code
```

The Activity Monitor displays the extended error code rather than the value 252 returned to the shell, in this case. For more information about error codes in Vault, see [“Errors Returned by the Vault Session”](#) on page 198.



The List of Images to be Vaulted

During a Vault session, Vault builds a list of images that are candidates for duplication or ejection. The `preview.list` file, which resides in the session directory for the current Vault session, includes all images that match the criteria specified on the profile **Choose Backups** tab except for the following:

- ◆ If a copy of an image already is in the Off-site Volume Group, that image will not be included in the `preview.list` file. Because images that have a copy in an Off-site Volume Group are already vaulted, Vault does not select them as candidates for vaulting.
- ◆ If the **Source Volume Group** field on the **Choose Backups** tab has been set to a specific volume group and if no copy of that image exists in that volume group, the image will not be added to the `preview.list` file.

After the `preview.list` file is generated, Vault evaluates the images in it to determine if they should be duplicated and/or ejected. Because several filters (that is, other profile configuration options) can exclude an image from duplication and ejection, the `preview.list` file may be a *superset* of the images that will eventually get duplicated by the session.

Duplication Exclusions

The following can eliminate an image from duplication:

- ◆ If **Disk Only** is specified on the **Choose Backups** tab, an image that has no disk copy will not be duplicated.
- ◆ If Vault determines that an image has already been duplicated, Vault will not duplicate the image again. Vault uses the following criteria to determine if an image has already been duplicated:
 - ◆ For One Copy Only. If the image exists in the Off-site Volume Pool, Vault does not duplicate it; conversely, if a copy of the image is *not* in the Off-site Volume Pool, Vault duplicates it.
 - ◆ For Concurrent Copies. Vault uses the **For Each Image If This Copy Fails** setting (**Continue** or **Fail All Copies**) to decide whether or not to duplicate an image. Each of the copies has its own **...If This Copy Fails** setting. Vault interprets the user's intent as follows:
 - ◆ **Continue**. If the setting for the copy is **Continue**, that copy is not critically important to the user. The duplication job will end with a partially successful (1) status if at least one of the other copies succeeds. If the current copy is the only one that fails, Vault will not re-duplicate the image the next time the profile runs. If all copies are set to **Continue**, at least one of those copies must exist or Vault will duplicate the image.

- ◆ **Fail All Copies.** If the setting for the copy is **Fail All Copies**, that copy is critically important to the user, and none of the copies will be successful. This forces Vault to retry the duplication the next time the profile runs if that image is selected for duplication (that is, if the time window of the profile allows that image to be selected again). However, if an unscheduled (and unlikely) event creates copies of the image, more than one copy of the image may be assigned to the destination volume pools. If the duplication operation results in more than the **Maximum Backup Copies**, the duplication step will fail. (**Maximum Backup Copies** are configured in **NetBackup Management > Host Properties > Master Server > *server_name* > Global NetBackup Attributes.**)

Ejection Exclusions

Vault ejects media listed in the `eject.list` file. If the profile skips the duplication step and an image in `preview.list` has no copy in an Off-site Volume Pool (configured on the **Eject** tab), it will not be ejected.

Vault Resiliency

The functionality that Vault uses to build the list of images to be duplicated and ejected allows Vault to do the following:

- ◆ Duplicate and/or eject images that were not processed because of a problem during the previous run of the profile. By configuring the image selection period to be a sufficient length of time, the Vault profile will duplicate an image if the duplication of that image failed during the previous run of that profile. To determine a sufficient length of time for the image selection period, see [“Overlap the Time Window in the Profile”](#) on page 22.
- ◆ Not duplicate images that were successfully duplicated by a previous job. You can restart a Vault session that was only partially successful, and Vault will not duplicate an image that was duplicated by a previous job. This provides for maximum flexibility or resiliency by allowing you to configure a longer image selection period without reduplicating images.

One Vault profile can vault original backup images from some NetBackup backup policies and also duplicate and vault images from other backup policies.



Ejecting Media

If you configure a profile to defer ejection, you must perform or schedule another action to eject media. (Conversely, if a profile that selects and/or duplicates images also ejects media, you do not have to perform another action to eject media.) You can use one of the following actions to eject media that was not ejected by a profile that selected or duplicated images:

- ◆ Manually by using the Vault Management node in the NetBackup Administration Console
- ◆ Manually by using the Vault Operator Menu
- ◆ Manually by using the `vlteject` command
- ◆ Automatically by creating and scheduling a Vault policy and entering the appropriate `vlteject` command and options in the file list

Note You must use one of the Vault methods to eject media; if you use a NetBackup or Media Manager option to eject media, the correct database entries will not be made and the Vault reports will not be accurate.

Related Topics

- ◆ [“Media Ejection Overview”](#) on page 88
- ◆ [“ACS MAP Overview”](#) on page 89
- ◆ [“Eject Mode \(Immediate or Deferred\)”](#) on page 89
- ◆ [“Media Ejection Timeout Impact”](#) on page 91

Previewing Media To Be Ejected

Before you eject media, you can preview the media that will be ejected. To preview that media, you can use the following:

▼ To use the Administration Console to preview media that will be ejected

1. Select the vault or profile for which you want to eject media.
2. Select **Actions > Deferred Eject**.

The Deferred Eject dialog appears. The options selected in the dialog depend on whether you are ejecting for all vaults, for a single vault, or for a profile.

3. If necessary, select a vault, profile, or session ID.

4. Click **Get Preview**, then select one or more of the profiles in the Eject Preview window.

▼ **To use the `vlteject` command to preview media that will be ejected**

1. From a command prompt, enter `vlteject` command in the following format, specifying the robot, vault, or session for which you want to preview ejected media:

```
vlteject -preview [-profile profile_name] [-robot robot_name] [-vault  
vault_name [-sessionid id]]
```

Ejecting Media by Using the NetBackup Administration Console

You can use the NetBackup Administration Console to eject media and generate reports for all vaults, for a single vault, or for a profile for which media have not yet been ejected.

When you select Deferred Eject, the default selections on the Deferred Eject dialog depend on whether you are ejecting for all vaults, for a single vault, or for a profile. From the dialog, you can initiate the eject operation or preview the media that will be ejected. The preview shows the session IDs for which the deferred eject will occur and the media IDs for each session selected. You also can select whether to generate the reports after the ejection.

▼ **To eject media by using the NetBackup Administration Console**

1. Select the vault or profile for which you want to eject media.
2. Select **Actions > Deferred Eject**.

The Deferred Eject dialog appears. The options selected in the dialog depend on whether you are ejecting for all vaults, for a single vault, or for a profile.

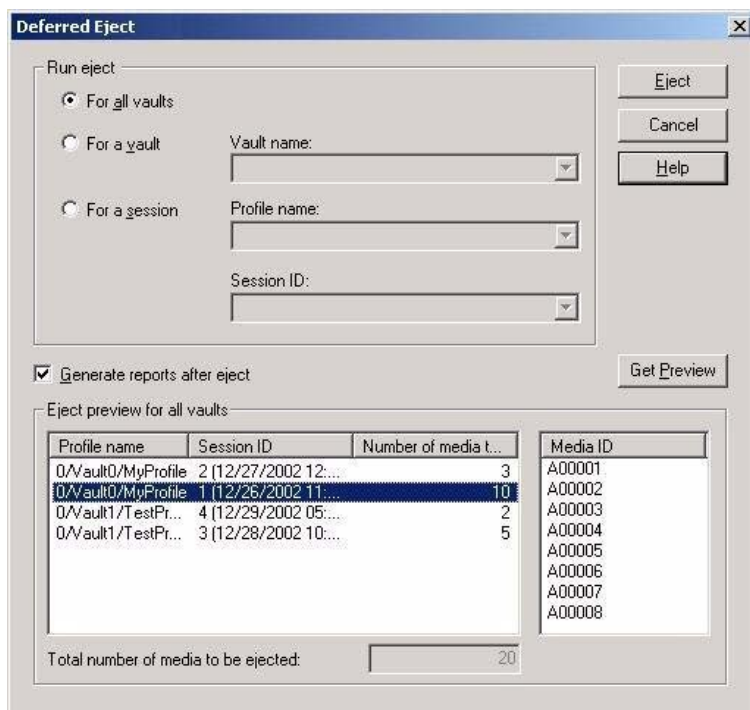
3. Select or change any of the options on the Deferred Eject dialog.
4. Click **Eject**.

To preview the media that will be ejected, click **Get Preview**, then select one or more of the profiles in the Eject Preview window.

To monitor the progress of or cancel the eject operation, use the NetBackup Administration Console Activity Monitor.



The following shows the Deferred Eject dialog with all vaults selected and previewing the media that will be ejected for the highlighted session:



Ejecting Media by Using the Vault Operator Menu

You can use the Vault Operator Menu to eject media and generate reports for Vault sessions for which media have not yet been ejected (the reports must be configured in the profiles). The Vault Operator Menu calls the `vlteject` command to accomplish the media ejection. You also can use the Vault Operator Menu to preview the media to be ejected.

For more information about using the Vault Operator Menu, see [“Using the Vault Operator Menu Interface”](#) on page 191.

▼ To eject media by using the Vault Operator Menu

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. If necessary, select a profile.

3. Select one of the following options:
 - ◆ Eject Media for This Session
 - ◆ Consolidate All Ejects
 - ◆ Consolidate All Reports and Ejects

Ejecting Media by Using the vlteject Command

You can use the `vlteject` command to eject media and generate reports for Vault sessions for which media have not yet been ejected (the reports must be configured in the profiles). The `vlteject` command can process the pending ejects and/or reports for all robots (that is, all sessions for all vaults), for all sessions for a single vault, or for a specific Vault session.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name] [-robot
robot_name] [-vault vault_name [-sessionid id]] [-auto y|n]
[-eject_delay seconds]
```

The `vlteject` command resides in the following directory:

UNIX: `/usr/opensv/netbackup/bin`
 Windows: `install_path\NetBackup\bin`

The following is a `vlteject` command example that ejects media for all robots that have sessions for which media has not yet been ejected and generates the reports:

```
vlteject -eject -report
```

The following example ejects all media that has not yet been ejected for all sessions for the CustomerDB vault and generates reports:

```
vlteject -vault CustomerDB -eject -report
```

The following is a `vlteject` command example that previews the media to be ejected for the CustomerDB vault:

```
vlteject -vault CustomerDB -preview
```

For more information about the `vlteject` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To eject media by using the vlteject command

1. In a terminal or command window, change to the directory in which the `vlteject` command resides.
2. Invoke the command, using the appropriate options and parameters.



Ejecting Media by Using a Vault Policy

You can use a Vault policy to eject media and/or generate reports for Vault sessions that have been completed already and for which media have not been ejected. In the Vault policy, specify Vault as the policy type, do not specify clients, and specify the `vlteject` command on the **Backup Selections** tab.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name] [-robot  
robot_name] [-vault vault_name [-sessionid id]] [-auto y|n]  
[-eject_delay seconds]
```

The `vlteject` command resides in the following directory:

UNIX: `/usr/opensv/netbackup/bin`

Windows: `install_path\NetBackup\bin`

The following is an `vlteject` command example that ejects media for all robots that have sessions for which media has not yet been ejected and generates the reports:

```
vlteject -eject -report
```

The following example ejects all media that has not yet been ejected for all sessions for the CustomerDB vault and generates reports:

```
vlteject -vault CustomerDB -eject -report
```

For more information about creating NetBackup policies, see the *NetBackup System Administrator's Guide, Volume I*. For more information about the `vlteject` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To create a Vault policy that ejects media

1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Click the **New Policy** button.
The Add a New Policy dialog appears.
3. Enter a unique name for the new policy in the Add a New Policy dialog.
4. Click **OK**.
The Add New Policy dialog appears.
5. On the **Attributes** tab, select **Vault** as the policy type.
6. On the **Schedules** tab, click **New** to create a new schedule.

The type of backup defaults to **Automatic Vault**.

7. Complete the schedule.
8. Bypass the **Client** tab (clients are not specified for Vault jobs).
9. On the **Backup Selections** tab, enter the `vlteject` command and the appropriate options for the policy.
10. Click **OK**.

Consolidating Ejects

You can eject media from more than one vault session, which is known as *consolidating* ejects. For example, you may use one vault policy to duplicate media daily but eject media only at the end of the week.

If you consolidate ejects, you should also consolidate reports. The individual reports are concatenated to produce the consolidated reports. If you consolidate your reports and also rename reports, you should use the same customized report title for all profiles whose reports will be consolidated.

Related Topics

- ◆ [“Configuring Eject”](#) on page 85
- ◆ [“Configuring Reports”](#) on page 91

▼ To consolidate ejects and reports for a profile

1. Select **Deferred Eject** on the profile **Eject** tab.
This action ensures that tapes will not be ejected automatically for each Vault session.
2. Select **Deferred Reports** on the profile **Reports** tab.
This action ensures that reports will not be generated automatically for each Vault session.
3. Eject media and generate reports using one of the methods in [“Ejecting Media”](#) on page 106.



Injecting Media

In a normal volume rotation, you have to inject media back into a robot after media expires and is returned from your off-site storage location so that it is available for reuse. You also may need to inject unexpired media for restore or disaster recovery operations.

Injecting media updates the NetBackup and Media Manager catalogs so that the correct location of the media is recorded. If the robot does not have a bar code reader to identify the media being injected, you still must use an inject option so the location of the media is updated in the databases. How you accomplish the process of injecting the media depends on the robot library:

- ◆ If your library has a media access port (MAP), you insert the media to be injected into the MAP and then use one of the injecting options discussed in this section to move that media from the MAP to the library slots. If the library has a bar code reader, the appropriate database changes are made automatically.
- ◆ If the library does not have a MAP, you insert the media into the library slots or into a cartridge which is then placed into the robot. If the library has a bar code reader, the appropriate database changes are made automatically.
- ◆ If your library does not have a bar code reader, you must use the Move media option of the NetBackup Administration Console so the databases are updated.

You can inject media as follows:

- ◆ [Injecting Media by Using the Administration Console](#)
- ◆ [Injecting Media by Using the Vault Operator Menu](#)
- ◆ [Injecting Media by Using the `vlthinject` Command](#)

On UNIX systems, you also can inject media by using the `vmadm` command. For information about injecting media and the inject functions available by robot, see the *NetBackup Media Manager System Administrator's Guide*.

Injecting Media by Using the Administration Console

Use the NetBackup Administration Console to inject media for libraries that have bar code readers and libraries that do not have bar code readers.

▼ To inject media by using the Administration Console

1. Insert the media into the robotic library slots or media access port.
2. In the NetBackup Administration Console, click **Media and Device Management > Media > Robots**.

3. Select the robotic library where you inserted the volume.
4. Click **Actions > Inventory Robot**.
5. In the Inventory operation section, select **Update volume configuration**.
6. If your robot has a media access port into which you placed the media, select **Empty media access port prior to update** in the Inventory operation section.
7. To configure advanced options, click **Advanced Options**.
8. To clear any previous display in the Results section, click **Clear Results**.
9. Click **Start** to start the update.
10. Repeat as necessary until all media are injected.

▼ **To inject media for robots without bar code readers**

1. Insert the media into the robotic library slots (or into the cartridge and then inject the cartridge into the robot).
2. In the NetBackup Administration Console, click **Media and Device Management > Media**.
3. Select the volume to be injected into the library.
4. Click **Actions > Move**.
5. In the Move Volumes dialog, select or enter the robot, volume group, and slot number. Use the First Slot Number field to enter the slot into which you placed the volume.
6. Click **OK** to move the volume.
7. Repeat as necessary until all media are injected.

For information about how to use the Administration Console to inject media and the inject functions available by robot, see the *NetBackup Media Manager System Administrator's Guide*.

Injecting Media by Using the Vault Operator Menu

You can use the Vault Operator Menu to inject media. The Vault Operator Menu calls the `vltnject` command to accomplish the media ejection.



▼ To inject media by using the Vault Operator Menu

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. If necessary, select a profile.
3. If necessary, select a session.
4. Load the media in the robot's media access port.
5. Select **Inject Media into Robot**.
6. Repeat until all media are injected into the robot.

For more information about using the Vault Operator Menu, see [“Using the Vault Operator Menu Interface”](#) on page 191.

Injecting Media by Using the `vltinject` Command

The `vltinject` command injects volumes into a robot and updates the Media Manager volume database. It requires as an option the name of a profile (if unique) or a robot number, vault, and profile name.

The following is the format of the `vltinject` command:

```
vltinject profile|robot/vault/profile
```

The following example command injects volumes that were vaulted by the Payroll profile and that have been returned from the off-site vault:

```
vltinject Payroll
```

The following example injects volumes that were vaulted by the Weekly profile in the Finance vault and that have been returned from the off-site vault:

```
vltinject 8/Finance/Weekly
```

For more information about the `vltinject` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To inject media by using the `vltinject` command

1. In a terminal or command window, change to the directory in which the `vltinject` command resides, as follows:

UNIX: `/usr/opensv/netbackup/bin`

Windows: `install_path\NetBackup\bin`

2. Load the media to be injected into the robot's media access port.
3. Invoke the command, using the appropriate options and parameters.
4. Repeat until all media are injected.

Vaulting and Managing Media in Containers

A container is a box in which you can place media and then transfer that box to your off-site storage location. When you configure a vault, you select whether the media are stored in containers or slots at your off-site storage location. Vault tracks, reports, and recalls your media regardless of how the media are transferred and stored off site.

After the media are ejected from your robot, you must add the media logically to containers by using either the Vault Operator Menu or the `vltcontainers` command. The options available for adding media to containers are as follows:

- ◆ Enter the container and media IDs by typing them in with the keyboard. Using this method, you can add media to more than one container.
- ◆ Scan the container and media IDs by using a keyboard interface bar code reader. (Keyboard interface readers are also known as keyboard *wedge* readers because they connect (or wedge) between the keyboard and the keyboard port on your computer.) Using this method, you can add media to more than one container.
- ◆ Read an input file that contains the IDs or numeric equivalents of bar codes of all the media that will be added to one container. If you have a bar code reader that can write to a file, you can scan media bar codes and use that output file as input for the `vltcontainers` command.
- ◆ Add all the media ejected by a specific session to one container.

The default return date of a container is the date of the volume in the container that will be returned the latest. You can change the return date during the container and media ID entry process or at any time thereafter before a container is recalled.

You also can delete a container from the NetBackup and Media Manager databases. If a container becomes empty after it is recalled and all media that reside in it are either injected back into the robot or assigned to another container, it is deleted from the NetBackup and Media Manager databases.

If you use containers, Vault reports on the containers and media outside the context of a profile or session.

To vault and manage containers and media, see the following:

- ◆ [Vaulting Media in Containers](#)
- ◆ [Managing Containers and Media](#)



- ◆ [Reporting on Containers and Media](#)

Vaulting Media in Containers

You can use either the Vault Operator Menu or the `vltcontainers` command to add media IDs to containers.

Vaulting Container Media by Using the Vault Operator Menu

After the media are ejected from your robot, you can use the Vault Operator Menu to enter the container and media IDs. For more information about using the Vault Operator Menu, see [“Using the Vault Operator Menu Interface”](#) on page 191.

▼ To vault media in containers by using the Vault Operator Menu

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. Eject the media that will be added to the containers.
3. Select **Container Management**.
4. Select one of the following options:
 - ◆ **Move media into one or more containers** if you intend to use the keyboard to enter media and container IDs or use a bar code scanner to scan the bar codes on the volumes and containers.
 - ◆ **Move all media ejected by this session, into one container** if you want to add all media ejected by a session to a container.
 - ◆ **Move all media listed in a file, into one container** if you want to add all media listed in an input file to a container.
5. Follow the prompts to complete the process of logically moving media into containers.

Vaulting Container Media by Using the `vltcontainers` Command

After the media are ejected from your robot, you can use the `vltcontainers` command to enter the container and media IDs. The following is the format of the `vltcontainers` command:

```
vltcontainers -run [-rn robot_number]
vltcontainers -run -usingbarcodes [-rn robot_number]
```

```
vltcontainers -run -vltcid container_id -vault vault_name -sessionid session_id
vltcontainers -run -vltcid container_id -f file_name [-rn robot_number] [-usingbarcodes]
vltcontainers -view [-vltcid container_id]
vltcontainers -change -vltcid container_id -rd return_date
vltcontainers -delete -vltcid container_id
vltcontainers -version
```

The following examples show how to use the `vltcontainers` command to add media to a container:

- ◆ To add the volumes ejected from robot number 0 to containers and enter the IDs by typing them in, use the following command:

```
vltcontainers -run -rn 0
```
- ◆ To add the volumes ejected from robot number 0 to containers and use a bar code reader to scan the container and media IDs, use the following command:

```
vltcontainers -run -usingbarcodes -rn 0
```
- ◆ To change the return date of container ABC123 to December 07, 2004, use the following command:

```
vltcontainers -change -vltcid ABC123 -rd 12/07/2004
```
- ◆ To delete container ABC123 from the NetBackup and Media Manager catalogs, use the following command:

```
vltcontainers -delete -vltcid ABC123
```
- ◆ To add all media ejected by session 4 of vault MyVault_Cntrs to container ABC123, use the following command:

```
vltcontainers -run -vltcid ABC123 -vault MyVault_Cntrs -sessionid 4
```
- ◆ To add media listed in file `C:\home\jack\medialist` that are ejected from robot number 0 to container ABC123, use the following command:

```
vltcontainers -run -vltcid ABC123 -f C:\home\jack\medialist -rn 0
```
- ◆ To add media to container ABC123 that was ejected from a robot that is attached to the master server and read the bar codes for that media from file `C:\home\jack\medialist`, use the following command:

```
vltcontainers -run -vltcid ABC123 -f C:\home\jack\medialist -usingbarcodes
```

For more information about the `vltcontainers` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.



▼ **To vault media in containers by using the `vltcontainers` command**

1. Eject the media that will be added to the containers.
2. Invoke the `vltcontainers` command, using the appropriate options and parameters.
3. Follow the prompts to move the media logically into containers.

Managing Containers and Media

After the media and containers are sent to your off-site storage location, you can still perform tasks to manage the containers and media. You can view and change return dates of containers. If a container has been recalled and is empty of media, you can delete the information about a container from the NetBackup and Media Manager databases.

Using the Vault Operator Menu to Manage Container Media

You can use the Vault Operator Menu to change a container return date and to delete the information about a container from the NetBackup and Media Manager databases. For more information about using the Vault Operator Menu, see “[Using the Vault Operator Menu Interface](#)” on page 191.

Note *Iron Mountain users:* To change a container return date, change the date using the Vault Operator Menu or the `vltcontainers` command then resend the Container Inventory Report or the Iron Mountain FTP file to Iron Mountain. Do not use the Iron Mountain account management facilities to change a container return date; if you do, the Vault reports will not match the report information maintained by Iron Mountain.

▼ **To view a container return date by using the Vault Operator Menu**

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. Select **Container Management > View a container’s return date.**
3. Follow the prompts to enter a container name.

▼ **To change a container return date by using the Vault Operator Menu**

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. Select **Container Management > Change a container’s return date.**

3. Follow the prompts to enter container names and change dates.

▼ **To delete a container by using the Vault Operator Menu**

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. Select **Container Management > Delete a container**.
3. Follow the prompts to enter a container name and delete the records of a container.

Note If a container becomes empty after it is recalled and all media that reside in it are either injected back into the robot or assigned to another container, it is deleted from the NetBackup and Media Manager databases.

Using the `vltcontainers` Command to Manage Container Media

You can use the `vltcontainers` command to view and change a container return date and to delete the information about a container from the NetBackup and Media Manager databases.

For more information about the `vltcontainers` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Note *Iron Mountain users:* To change a container return date, change the date using the Vault Operator Menu or the `vltcontainers` command then resend the Container Inventory Report or the Iron Mountain FTP file to Iron Mountain. Do not use the Iron Mountain account management facilities to change a container return date; if you do, the Vault reports will not match the report information maintained by Iron Mountain.

▼ **To view a container return date by using the `vltcontainers` command**

- ❖ Invoke the `vltcontainers` command using the `-view` option. For example, to view the return date of container ABC123, use the following command:

```
vltcontainers -view -vltcid ABC123
```

▼ **To change a container return date by using the `vltcontainers` command**

- ❖ Invoke the `vltcontainers` command using the `-change` option and specifying the `-vltcid` parameter and argument and `-rd` parameter and argument. For example, to change the return date of container ABC123 to December 07, 2004, use the following command:

```
vltcontainers -change -vltcid ABC123 -rd 12/07/2004
```



▼ **To delete a container by using the vltcontainers command**

- ❖ Invoke the `vltcontainers` command, using the `-delete` option and specifying the `-vltcid` parameter and argument. For example, to delete container ABC123 from the NetBackup and Media Manager catalogs, use the following command:

```
vltcontainers -delete -vltcid ABC123
```

To be deleted, a container must be empty.

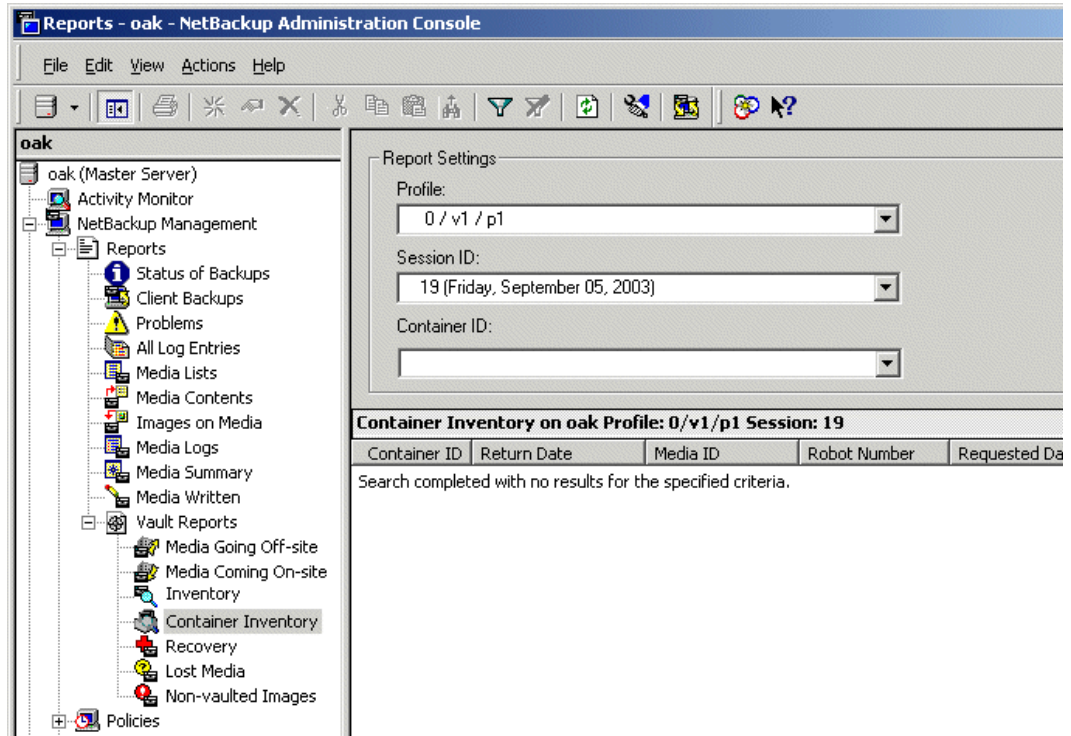
Reporting on Containers and Media

The Container Inventory Report shows all the containers configured in your vaulting environment, the return date of each container, and the media that are in each container. Alternatively, you can specify a container ID to generate a report of the media in a specific container.

If you are using containers, all of the other Vault reports will list the ID of the container in which the volume resides rather than a slot number. Reports will not show container information until after you add container and media IDs in Vault. Media are removed logically from a container when they are injected back into the robot.

You also can generate the Container Inventory Report by using the Vault Operator Menu (**Run Individual Reports > Container Inventory**). For more information, see [“Using the Vault Operator Menu Interface”](#) on page 191.

The following is an example of the Container Inventory Report window:



▼ To generate a Container Inventory report

1. In the NetBackup Administration Console, select **Reports > Vault Reports > Container Inventory**.
2. In the **Container ID** field, select All Containers or the ID of the container for which you want a report.
3. Click **Run Report**.

The details pane of the Administration Console will display the report details.

Assigning Multiple Retentions with One Profile

Different types of data often are handled differently. For example, you may want to vault your finance data for 7 years and your customer data for 20 years. To do this, the off-site copy of your backups will have different retentions based on the type of data. Vault can



process different types of data individually if your backups are organized based on the type of data being protected (for example, if you have separate backup policies based on the data type, such as a Finance backup policy and a CustomerDB backup policy).

When a Vault session creates one duplicate copy of many backups, it typically assigns the same retention to all of the duplicates created. Alternatively, you can configure Vault to calculate a retention for the duplicate copy based on the type of data. When doing duplication, there are three ways to handle retentions:

- ◆ Specifying **No Change** keeps the same expiration date as the original copy.
- ◆ Specifying a numeric retention level applies that retention, calculated from the backup time of the original image.
- ◆ Specifying **Use Mappings** instructs the Vault profile to find the correct retention for a specific type of data in a mapping file that you have configured.

The retention for a duplicate copy of a particular type of data is based on the retention level configured in the backup policy for that type of data. And, the retention level of the first (or only) copy configured in the schedule that created the image is used as the key into a mapping file. You configure the mapping file to convert the schedule's retention to a new retention for the duplicated copy.

For example, suppose you want to retain the on-site copy of all your data for 2 weeks, the off-site copy of your Finance data for 7 years, and the off-site copy of your CustomerDB data for 20 years. You can do this as follows:

1. Using Host Properties in the NetBackup Administration Console, configure retention levels 1 and 11 to be 2 weeks, retention level 12 to be 7 years, and retention level 13 to be 20 years.
2. In your Finance backup policy, assign retention level 1 (2 weeks) to the first (or only) copy configured in the schedule.
3. In your CustomerDB policy, assign retention level 11 (2 weeks) to the first (or only) copy configured in the schedule.
4. In your Vault profile, on the Duplication tab configure **Retention Level** to be **Use Mappings**.
5. Configure the mapping file as follows:

| | |
|---|----|
| 0 | 0 |
| 1 | 12 |
| 2 | 2 |
| 3 | 3 |

| | |
|----|----|
| 4 | 4 |
| 5 | 5 |
| 6 | 6 |
| 7 | 7 |
| 8 | 8 |
| 9 | 9 |
| 10 | 10 |
| 11 | 13 |
| 12 | 12 |

With this mapping, duplicated images from the Finance policy/schedule are assigned a retention level of 12 (7 years) and duplicated images from the CustomerDB policy/schedule are assigned a retention level of 13 (20 years). The duplicated images will be written to different media if the **Allow Multiple Retentions Per Media** property is not set (**NetBackup Management > Host Properties > Master Server > Media**).

A template `retention_mappings` file is provided on the NetBackup master server:

UNIX: `/usr/opensv/netbackup/db/vault/retention_mappings`

Windows: `install_path\NetBackup\db\vault\retention_mappings`

The mapping file is two columns by 25 lines (25 retention levels). The first column is the key into the retention mappings, and the second column specifies the retention level Vault will apply to the duplicated copy. (Again, given an image to duplicate, the retention level of the first (or only) copy configured in the schedule that created the image is used as the key into a mapping file.) By default, each retention level maps to itself (that is, retention level 0 maps to 0, 1 maps to 1, and so on). To change the retention mapping, change the value in column two. As in the above example, to retain the duplicated copy 7 years (retention level 12) when the retention in the policy/schedule is 2 weeks (retention level 1), specify "12" in the second column of the appropriate line in the mappings file.

You can configure the same retention periods for all vaults, separate retention periods for each vault, or any combination of vault specific and general mappings:

- ◆ To configure the same mappings for all vaults, configure the mappings in the template file.
- ◆ To use a separate mapping file for a specific vault, copy the template file and name it `retention_mappings.vault`, where *vault* is the name of the vault. Then, configure the mappings for that vault in that file. If a vault specific mapping file does not exist, a duplication rule configured to use a mappings file will use the `retention_mappings` file.



When you configure a profile, you can specify normal retention calculation for some duplication rules and alternative retention mappings for other duplication rules.

The retention period for the duplicate images begins on the date the primary backup image was created, not on the date the duplicate image was created.

If the backup policy that created the primary backup image no longer exists, the duplication of that image will fail and the job will continue but report partial success on completion.

The values for the retention levels are configured in NetBackup Host Properties > Master Server > Retention Levels. For information about configuring different values for the retention levels, see the *NetBackup System Administrator's Guide, Volume I*.

Vaulting Additional Volumes

Usually, you create the necessary copies of backup media during a NetBackup policy job or a Vault profile duplication job, and the Vault profile ejects the media for transfer off site. After the Vault profile is run, you cannot run the profile again to create additional copies of media that was already sent off site.

However, you can use other methods to create and eject additional copies of backup media after the NetBackup policy and Vault profile have been run. You can duplicate the volume manually or you can configure Vault to duplicate the volume. If you want to duplicate and eject one or several additional volumes one time only, the easiest solution is to duplicate the volume manually.

To duplicate an additional volume, the primary copy of the volume must be in the robot. If the primary copy is not in the robot but a duplicate copy is, you can use the `bpchangeprimary` command to change the duplicate to primary before you create an additional volume. For information about the `bpchangeprimary` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Duplicating a Volume Manually

If you use the following instructions to duplicate a volume manually, the tape will appear on the Picking List for Vault report when it expires and be recalled from the off-site vault as part of your normal operations.

▼ To duplicate a volume manually

1. Duplicate the volume manually by using the `bpduplicate` command. When duplicating the volume, specify the same off-site volume pool used for the volume already vaulted.
2. Assign the vault vendor's slot number for the volume by using the `vltoffsitemedia` command. The slot number is assigned in the `vltslot` field. You can assign values to other vault fields if desired.

Caution Do not assign a value to the `vlreturn` field. If you assign a value, the volume will never appear on the Picking List for Vault report.

3. Move the volume into the off-site volume group by using the `vmchange` command. Use the same off-site volume group as the first vaulted copy.

If the volume is in the same off-site volume group and the same off-site volume pool used by the regularly scheduled Vault profile, this volume will appear on the Picking List for Vault report when the first vaulted copy expires (if you did not assign a value to the `vlreturn` field).

4. Eject the volume.
5. Edit the file of the Picking List for Robot report to insert this volume into the list, then print the report and give it to the vault vendor.

For information about the `bpduplicate` and `vltoffsitemedia` commands, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Duplicating a Volume by Using Vault

If you use Vault to create and eject an additional copy of a volume that is already in the off-site vault, you must use a different vault, different off-site volume group, and different off-site volume pool than the first volume. The additional volume will not appear in the Picking List for Vault, so you must use the Lost Media Report to recall the volume after it expires.

▼ To duplicate a volume by using Vault

1. Create and configure a new vault. Use a different off-site volume group than the first volume.
2. Create and configure a new profile to duplicate and eject the volume. Assign the volume to a different off-site volume pool than the first volume.



3. Configure the eject step of the profile to search the off-site volume pool in which the additional volume was assigned.
4. Run the profile.
5. To recall the volume after it expires, run the Lost Media Report.

If you run the Lost Media Report as part of your normal operations, the volume will appear on the report after it expires.

Revaulting Unexpired Media

When you inject unexpired media from off-site storage back into a robot (for example, to perform a restore), you should revault the media. If you have to revault many tapes, you should create a new profile to revault them. If only a few tapes are to be revaulted, revaulting them manually may be the easiest and fastest option.

▼ To revault media by creating a new profile

1. Copy the original Vault profile that was used to eject the media.
2. In this new profile, change the Choose Backups time window to shift it far enough back in time so that it will select the images on the volumes that you want to revault.
3. Start a session using this new vault profile.

Vault will recognize that copies of images eligible to be vaulted exist and will not duplicate the images even if the duplication step is configured. The profile will eject the volumes to be revaulted.

4. If you are vaulting containers, logically add the volumes to containers. The container ID field is cleared when media vaulted in containers is injected back into the robot, so you must add the media to containers. See [“Vaulting and Managing Media in Containers”](#) on page 115).

If you are vaulting media in slots, Vault assumes that the media will be returned to the same slots in off-site storage from which they were recalled.

5. Delete the new profile you created to do the revaulting.
6. If you froze your media during the data restore process, use the `bpmedia` command to unfreeze it.

If you froze the media, you have to unfreeze the it so it will be recalled and returned to volume pool rotation when it expires. Vaulted media that are suspended will be unsuspended automatically when they expire and are recalled.

7. Return the media to your vault vendor so that all backups on that media will be available for future disaster recovery.
8. Run the Recovery Report to ensure that media are available in off-site storage for future use.

For information about the `bpmedia` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To revault media manually

1. Manually eject the media, using one of the following methods:
 - Use the `vmchange` command.
 - In the NetBackup Administration Console, highlight the media ID then select the **Eject Volumes from Robot....** operation on the **Actions** menu.

Note `vlteject` and `vltopmenu` will not work for this purpose.

2. Manually transfer the media to the off-site volume group, using one of the following methods:
 - Use the `vmchange` command.
 - In the NetBackup Administration Console, highlight the media ID then select the **Change Volume Group....** operation on the **Actions** menu.
3. If you are vaulting containers, logically add the volumes to containers. The container ID field is cleared when media vaulted in containers is injected back into the robot, so you must add the media to containers. See [“Vaulting and Managing Media in Containers”](#) on page 115).

If you are vaulting media in slots, Vault assumes that the media will be returned to the same slots in off-site storage from which they were recalled.

4. If you froze your media during the data restore process, use the `bpmedia` command to unfreeze it.

If you froze the media, you have to unfreeze the it so it will be recalled and returned to volume pool rotation when it expires. Vaulted media that are suspended will be unsuspended automatically when they expire and are recalled.



5. Return the media to your vault vendor so that all backups on that media will be available for future disaster recovery.
6. Run the Recovery report to ensure that the media are available for future disaster recovery operations.

For information about the `vmchange` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Tracking Volumes Not Ejected by Vault

Eject is the event that Vault uses to update the NetBackup database for location and to recall volumes. In normal operation, Vault must eject volumes so they are tracked at off-site storage and recalled after they expire.

If you ejected volumes or removed them from your robot without using a Vault eject method, it is still possible to use Vault to track them — if they are in a volume pool you are using for off-site media. Using NetBackup commands, you can change the attributes for volumes that were not ejected by Vault so they will appear in the Vault reports and be recalled when they expire.

Note This process will work *only* if the volumes already are in a volume pool on the Eject tab of a profile in the vault you specify in the command. You cannot change the pool of an assigned volume.

First, use the `vmchange` command to change the volume group of the volumes, which Vault uses to track their location. For example, the following `vmchange` command changes the volume group of volume A00001:

```
vmchange -new_v offsite_volgrp -m A00001
```

Then, use the `vltoffsitemedia` command to change the Vault specific attributes. The following `vltoffsitemedia` example changes the vault attributes of volume A00001:

```
vltoffsitemedia -change -m A00001 -vltname offsite_vault -vltsent  
07/03/2004 -vltreturn 0 -vltslot 99 -vltsession 33
```

If you are adding the volumes to slots at your off-site storage location, use the `vltoffsitemedia` command with the `-list` option to find empty slots into which you can add the volumes. For a catalog volume return date, see the table on the following page.

If you are placing the volumes in containers, use the `vltcontainers` command to add the volumes logically to containers after you specify the off-site volume group and the vault attributes (see [“Vaulting and Managing Media in Containers”](#) on page 115). The default return date of a container is the date of the volume in the container that will be returned the latest; you may have to change the container return date if the volume you are adding expires later than any volume already in the container.

The following are the `vltoffline` options you can use to set the necessary volume attributes:

| Option | Description |
|-------------------------------------|--|
| <code>-vltname vault_name</code> | The name of the vault. |
| <code>-vltreturn date</code> | Set the return date to 0; Vault uses the latest expiration date of the images on the volume as the return date. <i>Exception:</i> if the volume is a NetBackup catalog backup volume, set the date the volume should be returned from offsite. |
| <code>-vltsent date</code> | Set the sent date to the date the volume was ejected. The format of the date depends on your locale setting. For the C locale, the date syntax is mm/ dd/ yyyy [hh[:mm[: ss]]]. |
| <code>-vltsession session_id</code> | The ID of the session that vaulted the volume or container. Set it to a nonzero number that is different from existing session IDs |
| <code>-vltslot slot_id</code> | The ID of the slot in which the volume resides in the off-site vault. Ensure that this is an empty slot at your off-site storage location. If you are placing the volume in a container, do not specify this option. |

Vaulting Media Not Created by NetBackup

Vault can eject and track media that was not created by NetBackup if the media are managed by Media Manager. Vault executes shell scripts, called *notify scripts*, at specific checkpoints during the vault process. If you place an eject notify script in the NetBackup `bin` directory, Vault executes that script immediately before it performs the eject step of a profile. If that script includes valid Media Manager media IDs, Vault ejects that media along with the media selected by the Vault profile *if* the volume pool in which that media resides is in the Off-site Volume Pools list on the **Eject** tab of the Profile dialog. Vault will eject notify script media even if no other media are selected for ejection by the Vault profile.

When expired media are injected back into the robot, all Vault specific fields for that media in the volume database are cleared. The media description field also is cleared if the `VAULT_CLEAR_MEDIA_DESC` parameter is set in the `vm.conf` file (see [“Clearing the Media Description Field”](#) on page 134).

Templates of the various notify scripts are provided with Vault. You can copy and modify the `vlt_ejectlist_notify` script so you can eject the non NetBackup media. For information about how to use the notify scripts, see [“Using Notify Scripts”](#) on page 131. The scripts include information about how to modify and test them.



▼ **To vault non NetBackup media managed by Media Manager**

1. Copy the `vlt_ejectlist_notify` script and name it appropriately (that is, add the appropriate extension to the name).
2. Edit the script as follows:
 - a. Add the media IDs of the non NetBackup media that you want to eject.
 - b. Add an expiration date to the media by using the `vltoffsetmedia` command with the `vltreturn` option.

The script will execute the `vltoffsetmedia` command(s) and assign the expiration date(s). The media will appear on the Picking List for Vault on the date it expires.

3. Place the script in the NetBackup `bin` directory.
4. Configure a Vault profile so that it includes the volume pool in which the media are assigned in the Off-site Volume Pools list on the **Eject** tab of the Profile dialog.
5. Start the Vault profile that includes the volume pool in which the media are assigned in the Off-site Volume Pools list on the **Eject** tab of the Profile dialog.

If the script executes successfully, the media will be ejected. Vault will assign vendor slot or container IDs to the volumes, and they will appear on the Picking List for Robot report.

For information about the `vltoffsetmedia` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Notifying a Tape Operator When Eject Begins

You can use the `vlt_starteject_notify` script to notify a tape operator when the eject process begins. For more information about notify scripts, see [“Using Notify Scripts”](#) on page 131. The scripts include information about how to modify and test them.

Before you use a notify script, ensure that your systems are set up properly for e-mail. See [“Setting Up E-Mail”](#) on page 179.

▼ **To set up eject notification for tape operators**

1. Copy the `vlt_starteject_notify` script and name it appropriately (that is, add the appropriate extension (robot, vault, and/or profile) to the name).

On Windows systems, notify scripts have a `.cmd` extension.



2. Edit the line in the script that includes the "Vault start eject notification" subject as follows:
 - a. Remove the comment characters at the beginning of the line.
 - b. Replace the `$someone_who_cares` (UNIX) or `someone_who_cares` (Windows) string with the e-mail address of the person to be notified or the e-mail alias for the people to be notified.
3. Make any other changes in the script that are necessary.
4. Place the script in the NetBackup `bin` directory.

Using Notify Scripts

A Vault job can call *notify scripts* at specific points during the execution of the job. Vault includes template notify scripts that you can customize for your use. You can use a script for a robot, a vault, or a profile. The template notify scripts are in the following directory:

UNIX: `/usr/opensv/netbackup/bin/goodies`

Windows: `install_path\NetBackup\bin\goodies`

On Windows systems, the names of the scripts include a `.cmd` extension. They include instructions that can help you edit them for your needs.

To be called and executed, a script must be copied to the NetBackup `bin` directory. A script must return a normal status (0) for the Vault job to continue processing. In case of failure, the script must return a nonzero status code to cause the Vault job to stop. On UNIX systems, the return status is communicated to the Vault job through the exit call. On Windows systems, the scripts communicate the return status in a file defined by the `EXIT_STATUS` environment variable, which is set by Vault.

The following scripts have been provided with Vault:

`vlt_start_notify`

Called by the Vault session after it starts. For example, you can use it to send notification when the Vault job begins.



| | |
|------------------------------------|--|
| <code>vlt_ejectlist_notify</code> | Called by the Vault session before the list of media to be ejected (the <code>eject.list</code>) is built. You can use this script to add media managed by Media Manager but not created by NetBackup or Vault to the eject list. The script writes media IDs to the <code>addon_medialist</code> file; Vault reads the <code>addon_medialist</code> file and ejects the media listed in that file during the current Vault session <i>if</i> the volume pool in which that media resides is in the Off-site Volume Pools list on the Eject tab of the Profile dialog. |
| <code>vlt_starteject_notify</code> | Called by the Vault session after the <code>eject.list</code> file is built and before the automatic eject process begins. You can use this script to send notification when the eject process begins or perhaps to suspend the media in the eject list. If the eject step is not configured for the profile, the <code>vlt_starteject_notify</code> script is not called. |
| <code>vlt_endeject_notify</code> | Called at the end of eject processing. You can use this script to send notification when the eject process ends. If the eject step is not configured for the profile, the <code>vlt_endeject_notify</code> script is not called. |
| <code>vlt_end_notify</code> | Called by the Vault session immediately before it exits. One use for this script is to start another vault job; you can then run Vault jobs in succession and avoid resource contention. |

Before you use a notify script, ensure that your systems are set up properly for e-mail. See [“Setting Up E-Mail”](#) on page 179.

Notify Script for a Specific Robot

You can use a notify script to create a unique, customized script for each robot in your configuration. To create a notify script for a specific robot, append the robot number to the script name and copy the script to the NetBackup bin directory.

For example, a `vlt_start_notify` script for a specific robot will appear as follows:

```
vlt_start_notify.robot_number
```

The script will be executed for all profiles created for the robot.

Use the same methodology to create other notify scripts.

Notify Script for a Specific Vault

You can use a notify script to create a unique, customized script for each vault in your configuration. To create a notify script for a specific vault, append the robot number and vault name to the script name and copy the script to the NetBackup `bin` directory.

For example, a `vlt_start_notify` script for a specific robot/vault combination will appear as follows:

```
vlt_start_notify.robot_number.vault_name
```

The script will be executed for all profiles created for a specific vault.

Use the same methodology to create other notify scripts.

Notify Script for a Specific Profile

You can use a notify script to create a unique, customized script for each profile in your configuration. To create a notify script for a specific profile, append the robot number, vault name, and profile name to the script name and copy the script to the NetBackup `bin` directory.

For example, a `vlt_start_notify` script for a specific robot/vault/profile combination will appear as follows:

```
vlt_start_notify.robot_number.vault_name.profile_name
```

This script will be executed for a specific profile defined for a specific vault.

Use the same methodology to create other notify scripts.

Order of Execution

The notify scripts are executed in specific to general order, as follows:

1. *script_name.robot_number.vault_name.profile_name*
2. *script_name.robot_number.vault_name*
3. *script_name.robot_number*
4. *script_name*



Clearing the Media Description Field

When media are returned from the off-site vault during a typical volume rotation, they are expired and ready for reuse. To avoid confusion, it may be helpful to clear the media description information when an expired volume is returned to the robot.

You can configure NetBackup so that the media description field is cleared when volumes are returned to the robot. To do so, set the `VAULT_CLEAR_MEDIA_DESC` parameter in the `vm.conf` file. The media description field will be cleared when other Vault information is cleared from the Media Manager volume database; the vault fields are cleared when the media are:

- ◆ Unassigned while in a robotic volume group
- ◆ Moved into a robotic volume group and then unassigned

Volume pool, volume group, and media description fields are used for all volumes, not just volumes used by Vault. The following are the Media Manager database fields dedicated to Vault information:

| Field | Description |
|-----------------------|--|
| <code>vltcid</code> | The ID of the container (container vaulting only). |
| <code>vltname</code> | The name of the vault. |
| <code>vlreturn</code> | The date the volume or container should be returned from the off-site vault. |
| <code>vltsent</code> | The date the volume or container was sent off-site. |
| <code>vltsid</code> | The ID of the session that vaulted the volume or container. |
| <code>vltslot</code> | The ID of the slot in which the volume resides in the off-site vault (slot vaulting only). |

For information about setting the `VAULT_CLEAR_MEDIA_DESC` parameter in the `vm.conf` file, see the *NetBackup Media Manager System Administrator's Guide*.

Ensuring Available Media for Catalog Backups

The Catalog Backup operation is different from a typical backup operation, as follows:

- ◆ Catalog backup tapes do not expire as regular backup tapes do. Vault recalls and reuses catalog backup media after the **Retention Period** configured for the catalog backup has passed. Vault unassigns any catalog backup media that appears on the Picking List for Vault or Distribution List for Robot and returns that media to the catalog volume pool so it is available to reuse as catalog backup media.
- ◆ Catalog backup media that were allocated from a scratch pool are not returned to the scratch pool. If NetBackup tries to write normal backup data to a volume that has catalog header information, including a volume that is an unassigned vaulted catalog tape, NetBackup will freeze that tape. Therefore, if they were returned to the scratch pool they would be frozen and unavailable for use.

If catalog backup media does not appear on the Picking List for Vault or Distribution List for Robot, it will not be unassigned. That media will appear on the Lost Media Report; you must unassign that media manually so it will be available for reuse (see [“Deassigning Vaulted NetBackup Catalog Media”](#) on page 135).

If you want to use a catalog volume as a general NetBackup backup tape, you must also relabel the tape by using the `bplabel` command. For information about `bplabel`, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

Note If your catalog backup spans volumes (multitape catalog backup method), you must not use the volume pool used for the regular catalog backup operation. Rather, it must use a volume pool included in the volume pool list for the eject process. These tapes are handled just like any other tape used by a backup policy and should use a volume pool used by other backup policies.

Deassigning Vaulted NetBackup Catalog Media

Usually, NetBackup catalog media are returned for use as catalog backup media after the retention period for the media has passed. If you want to reuse vaulted catalog backup media before it expires or if catalog backup media was not recalled properly from off-site storage, you must deassign that media manually by using the `vmquery` command.

After you deassign the media, you can use it again for a catalog backup. If you want to use the tape as a general NetBackup backup tape, you must also relabel the tape by using `bplabel`. For information about `bplabel` and `vmquery`, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

The `bplabel` command is in the following directory:

UNIX: `/usr/opensv/netbackup/bin/admincmd`

Windows: `install_path\NetBackup\bin\admincmd`

The `vmquery` command is in the following directory:



UNIX: `/usr/opensv/netbackup/volmgr/bin`

Windows: `install_path\NetBackup\Volmgr\bin`

▼ To deassign a catalog volume

1. Obtain the pool number for the volume by using the `vmquery` command, as shown in the following example for volume S04440:

```
vmquery -m S04440
=====
media ID:S04440
media type:8MM cartridge tape (4)
barcode:-----
description:CH_V1|101|S278|00000000
volume pool:Offsite_Catalog (3)
robot type:NONE - Not Robotic (0)
volume group:vault_grp
created:Tue Sep 3 10:08:32 2000
assigned:Tue May 6 00:11:45 2001
last mounted:Tue May 6 11:34:25 2001
first mount:Tue Sep 3 18:20:48 2000
expiration date:---
number of mounts:21
max mounts allowed:---
status:0x1
=====
```

The pool number is listed on the `volume pool` line, and is the number in between the parentheses. In this case, the media ID would be S04440 and the pool number would be 3.

2. Deassign the volume by using the `vmquery` command, as in the example below:

```
vmquery -deassignbyid S04440 3 1
```

Restoring Data from Vaulted Media

You may need to restore images from media that is stored in your off-site vault. The high-level procedure in this section describes how to restore data from vaulted media.



▼ **To restore data from vaulted media**

1. Recall the media.
2. Change the images to be recovered to primary (NetBackup restores from the primary image).

Use the `bpchangeprimary` command to promote a copy to primary. For information about the `bpchangeprimary` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* guide.
3. If the media is not suspended or frozen, suspend the media.

Use the `bpmedia` command to suspend the media. For information about the `bpmedia` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* guide.
4. Inject the media into the robot. For procedures, see [“Injecting Media”](#) on page 112.

Injecting the media moves it into the robot and also changes the off-site volume group attribute of the media to the robotic volume group so NetBackup knows that the media are in the robot.
5. Using the Backup, Archive, and Restore interface, restore the data. For procedures, see the *NetBackup User’s Guide for UNIX* or *NetBackup User’s Guide for Microsoft Windows*.
6. After restoring all the data, revault the media. For procedures, see [“Revaulting Unexpired Media”](#) on page 126.

Replacing Damaged Media

If media in your robot is damaged, you can use copies of the media (if available) from your off-site storage location to replace the damaged media. You also can use this process to recover images if the primary backup has expired, the volume has been overwritten and a copy in off-site storage is still available.

Note This image recovery process assumes that the NetBackup system and image catalog are current and up-to-date.

The instructions use an example to illustrate how to invoke the various commands used in the recovery process. Modify the command examples as appropriate for your purposes. Most of the commands used to recover from damaged media are in the following directory:



UNIX: `/usr/opensv/netbackup/bin/admincmd`

Windows: `install_path\netbackup\bin\admincmd`

After you recover and restore the damaged media, you should revault the media so that it again is available for recovery. For revaulting procedures, see [“Revaulting Unexpired Media”](#) on page 126.

▼ To replace damaged media

1. Identify the damaged media.

When you receive an error message during a restore, the errors are logged to the restore log and also show up on the Activity Monitor as the restore fails. You can set up a procedure using NetBackup scripts to send errors to an event management console to notify the storage administrator immediately of this type of media error.

2. Determine which backup images were on the damaged tape.

All images on a specific tape can be identified by running the `bpimmedia` command. It scans the entire NetBackup image catalog, so it may take a few minutes depending on the size of that catalog. For example, the following shows that volume S05423 contains one image from client `fgolddust`. It also shows that this image has been duplicated because it has (FRAG 2) entries. The full image name is `“fgolddust_0862806643”`:

```
# bpimmedia -mediaid S05423
```

```
IMAGE fgolddust 2 fgolddust_0862806643 golddust_BR1 0 Full_Weekly 0 3 19360
8654 85043 0 0
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0 *NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0 *NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0 *NULL*
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0 *NULL*
```

3. Determine which duplicate tapes were used and their host.

In step 2, the (FRAG 2) entries show that an image has been duplicated: the (FRAG 2 1) entry is the duplicate copy. On copy 1 there were 4 fragments (usually due to multiplexing). The (FRAG 2 -1) entry is the true image restore duplicate. In this case, the image `fgolddust_0862806643` is using media `S04440` for duplicating all of the original fragments. This is normal because the original image was multiplexed onto 4 tapes, while the duplicate was de-multiplexed during image duplication, and could fit on one tape.

Also note that the host for the media is printed for each fragment, in this case nirvana. With media servers, the host could be different than the master server. Under Vault, the duplication should normally occur on the same server that made the original backup, so the host server names would be the same for both copies of the image.

You can confirm this information by using `bpimagelist` command, as follows:

```
# bpimagelist -backupid fgolddust_0862806643

IMAGE fgolddust 0 0 2 fgolddust_0862806643 golddust_BR1 0 *NULL* root
Full_Weekly 0 3 862806643 4591 865485043 0 0 2356562 19360 2 7 1
golddust_BR1_0862806643_FULL.f *NULL* *NULL* 0 1 0 2 865830643 *NULL* 1 0 0
0 0 *NULL*
HISTO -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0 *NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0 *NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0 *NULL*
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0 *NULL*
```

To confirm which is the primary copy (the copy to be used for restores), use the `-L` option with `bpimagelist`, as follows:

```
UNIX: # bpimagelist -L -backupid fgolddust_0862806643 | grep Primary
Primary Copy: 1
Windows: bpimagelist -L -backupid fgolddust_0862806643 | find Primary
Primary Copy: 1
```

4. Tell NetBackup to use the duplicated copy rather than the original.

Execute the `bpimage -npc` command and option to change the primary copy. The new primary copy is used for restoring an image:

```
# bpchangeprimary -copy 2 -id fgolddust_0862806643 -cl fgolddust
```

To confirm the new primary copy, use the following command:

```
UNIX: # bpimagelist -L -backupid fgolddust_0862806643 | grep Primary
Primary Copy: 2
Windows: bpimagelist -L -backupid fgolddust_0862806643 | find "Primary"
Primary Copy: 2
```

5. Freeze the duplicated copy to ensure restore.

Use the command `bpmedia -freeze` to prevent NetBackup from expiring the images on the media and to ensure the media is assigned in Media Manager. You should also use the media host for this image, which was printed by `bpimmedia` in step 2. This is required when the host is different than the machine on which you are running this command.



```
bpmedia -freeze -m S04440 -host nirvana
```

6. Recall media from the vault.

Recall the appropriate volume from off-site storage. To determine the media ID, slot number, or container ID of the tape to recall, you can use the `vmquery` command, located in the following directory:

UNIX: `/usr/opensv/volmgr/bin`
Windows: `install_path\volmgr\bin`.

In the following example, the slot number (S278) is listed in the vault slot field:

```
vmquery -m S04440
=====
media ID:                S04440
media type:              8MM cartridge tape (4)
barcode:                S04440
media description:       Added by Media Manager
volume pool:             Vaulted_CustomerDB (2)
robot type:              NONE - Not Robotic (0)
volume group:            DB_offsite_volumes
vault name:              Customer_DB_Vault
vault sent date:         ---
vault return date:       ---
vault slot:              S278
vault session id:        1
created:                 Tue Sep 3 10:08:32 2000
assigned:                 Tue May 6 00:11:45 2001
last mounted:            Tue May 6 11:34:25 2001
first mount:             Tue Sep 3 18:20:48 2000
expiration date:         ---
number of mounts:        21
max mounts allowed:      ---
=====
```

7. Inject recalled media back into the robot.

When the tape is returned from the off-site vendor, inject it into the appropriate robotic library. First Insert the tape into the robot media access port. Then, from the NetBackup Administration Console, choose **Media and Devices Management**. Choose the **Inventory Robot...** option. Select the **Empty Media Access Port Prior to Update** checkbox.

You can also perform this function using the `vltinject` command.

8. Perform a normal restore.

Perform a normal restore operation. The restore should read from the new primary copy. The restore log should show a mount request for the duplicate media.

9. Unfreeze media used for duplicates.

After the restore is successful, unfreeze the duplicate media to allow the normal expiration process to be followed. If you want to send the tape off-site again, either remove it from the robot or leave it in the robot as the primary copy. VERITAS recommends that you suspend the media so that no images are written to it.

```
bpmedia -unfreeze -m S04440 -host nirvana
```

10. Create new duplicate images.

Optionally, you can create new duplicate images for transfer to your off-site vault vendor. For more information, see [“Bad or Missing Duplicate Tape”](#) on page 199.

11. Modify the NetBackup catalog for a large number of images.

In a disaster recovery situation in which a large number of images need to have their primary copy modified, run the `bpchangeprimary` command. This command will change the primary copy of all the backup images in the off-site volume pool for which the media was returned from the off-site vault.





Creating Originals or Copies Concurrently

You can create multiple copies of a backup image concurrently. Those copies are created concurrently by the Inline Tape Copy feature. For more information, see the following:

- ◆ [“Understanding Concurrent Copies”](#) on page 143
- ◆ [“Continue or Fail for Concurrent Copies”](#) on page 144
- ◆ [“Creating Original Images Concurrently”](#) on page 145
- ◆ [“Creating Duplicate Images Concurrently”](#) on page 147

Understanding Concurrent Copies

You can create up to four copies of the same backup image concurrently. If the images are created during a NetBackup policy job, all are considered original images. If the images are duplicated by using the NetBackup Administration Console Catalog node or during a Vault job, they are considered duplicate images.

NetBackup must be configured to allow a sufficient number of copies in the **Maximum Backup Copies** field for the NetBackup master server. (Configured in **NetBackup Management > Host Properties > Master Server > *server_name* > Global NetBackup Attributes**.) By default, the value is two.

All storage units must be connected to the same media server. Also, the storage unit must be configured to allow a sufficient number of concurrent jobs to support the concurrent copies (Maximum Concurrent Jobs or Maximum Concurrent Drives Used for Backup setting). You can write images concurrently to the following storage units:

- ◆ Media Manager storage units. If the Media Manager storage unit has more than one drive, the source and destination storage units can be the same.
- ◆ Disk storage units.
- ◆ Disk staging storage units.
- ◆ Network Data Management Protocol (NDMP) storage units *only* during Vault duplication and only one copy is allowed per duplication rule (NDMP is not supported during original backup). If the NDMP storage unit has more than one drive, the source and destination storage units can be the same.



Although specifying an NDMP storage unit restricts the number of copies to one, you can use multiple duplication rules to specify other storage units for images created by other media servers. For example, you can use one duplication rule to read an image from one media server and write a copy to an NDMP storage unit and use another duplication rule to read an image from a different media server and write copies to other storage units. (To specify multiple duplication rules in a Vault profile, select **Advanced Configuration** on the Profile dialog **Duplication** tab.) Because of potential NDMP performance limitations, VERITAS suggests that you duplicate between disk and tape drives that are directly attached to the same NDMP host.

If you create multiple original images concurrently during a NetBackup policy job, the backup time required may be longer than for one copy. Also, if you specify both Media Manager and disk storage units, the duration of disk write operations will match that of slower removable media write operations.

You cannot create images concurrently using the following:

- ◆ Storage unit groups
- ◆ Optical devices
- ◆ Quarter-inch cartridge (QIC) devices
- ◆ Third-party copies

Continue or Fail for Concurrent Copies

When making multiple copies of images concurrently, you can choose how an operation will behave if one of the copies fails. Your choice also can determine whether copies will be ejected, depending on the success or failure of the copy operation. It is possible for a duplication operation to succeed but no ejection to occur.

In NetBackup, your continue or fail choice affects only the current image copy; in Vault, your choice affects all copies of that image. By default, the option is configured to **Continue** in NetBackup and to **Fail All Copies** in Vault.

Continue Copies

If you choose **Continue** for all copies, the concurrent copy job is considered successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted; it is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool for ejection.

To ensure that media are ejected even if a concurrent copy operation fails during a NetBackup policy backup, do one of the following:

- ◆ Configure a Vault profile to duplicate the image, assign the copy to the off-site volume pool, and select **Fail All Copies**. If the copy fails during the original NetBackup backup job, the Vault profile will subsequently duplicate it. If the copy succeeds during the original backup job, the Vault profile will not duplicate it. Either way, a copy will be ejected for transfer off site.
- ◆ Monitor the Activity Monitor for a failed status for the copy that is assigned to the off-site volume pool. If that copy fails, duplicate that image and assign it to the off-site volume pool so it will be ejected. You can use the Administration Console Catalog node or the `bpduplicate` command to duplicate the copy.

Fail All Copies

The behavior of the fail option and the default settings depend on whether the concurrent copies operation was configured in Vault or in NetBackup:

- ◆ In Vault, if you choose **Fail All Copies**, all copies *of that image* will fail, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault will again try to duplicate the image if the following conditions are true:
 - ◆ The image is selected.
 - ◆ The Vault profile did not eject the primary backup.
- ◆ In NetBackup, if you choose **Fail All Copies**, the entire backup job will fail and no copies will be made. In this case, normal NetBackup behavior will ensure that a successful backup for this policy eventually occurs. That is, NetBackup will automatically retry the backup if time permits and, the next time the backup window for the policy opens, NetBackup will again try to run the backup (regardless of the frequency of the schedule). NetBackup will do this until the backup succeeds, although one or more backup windows may pass before the backup is successful.

Creating Original Images Concurrently

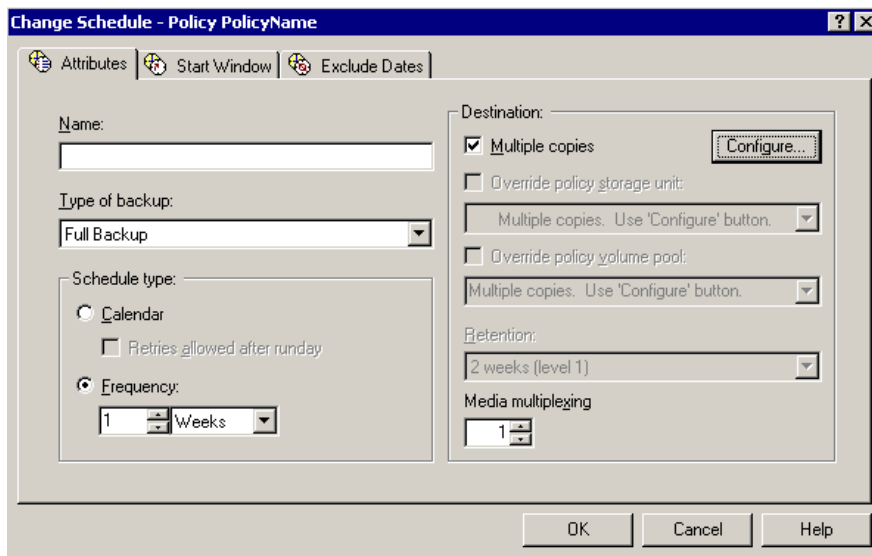
In a NetBackup policy job, you can create multiple original backup images concurrently. Vaulting original images has many benefits, including easier configuration of Vault, fewer chances for resource contention, and possibly fewer drives required.

▼ To create multiple backup images concurrently through the Policy node

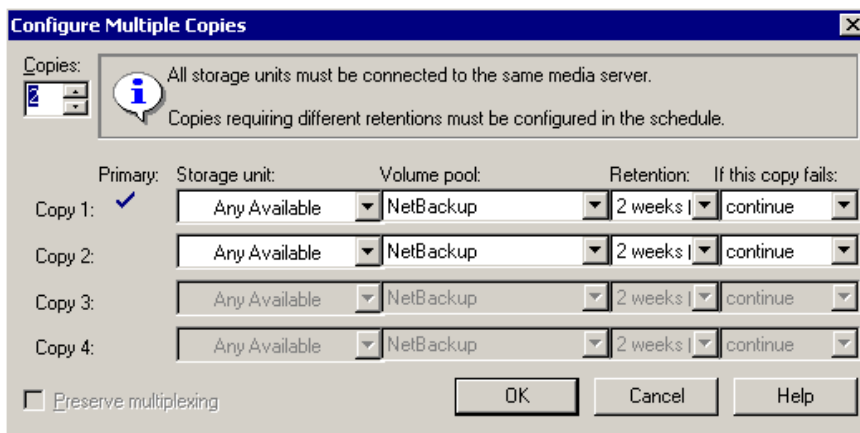
1. In the NetBackup Administration Console, expand **Master Server > NetBackup Management > Policies**.
2. Double-click an existing policy or:



- ◆ Windows: select **Actions > New > New Policy**
 - ◆ UNIX: click **Add New Policy** to create a new policy.
3. Select the Schedules tab.
 4. Double-click an existing schedule or click **New** to create a new schedule.
The Schedule dialog appears.



5. In the Schedule Attributes tab, select **Multiple Copies**, then click **Configure**.
The Configure Multiple Copies dialog appears.



6. Specify the number of copies to be created simultaneously.

The maximum is four. Copy 1 is the primary copy. If copy 1 fails, the first successful copy is the primary copy.

7. Specify the storage unit where each copy will be stored.

If a Media Manager storage unit has more than one drive, it can be used for both the source and the destination. Network Data Management Protocol (NDMP) storage units are not supported when creating multiple copies during a NetBackup policy job.

8. Specify the volume pool to which each copy will be assigned.

9. Select the retention level for each copy.

If you select **No Change**, the expiration date will be the same for the duplicate and original copies.

If you select a different retention period, the expiration date of the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2003, and its retention period is one week, the new copy's expiration date is November 21, 2003.

10. Select whether to **Continue** the other copies if a copy operation fails or to **Fail All Copies**.

11. Click **OK**.

12. Configure other schedule criteria as appropriate.

Creating Duplicate Images Concurrently

You can create multiple duplicate backup images concurrently either by using the NetBackup Catalog node or by configuring the Duplication tab of a Vault profile. Duplication is not always possible, so you must understand when you can use duplication in NetBackup.



When Duplication is Possible

The following describes when duplication is and is not possible in NetBackup:

| Possible to duplicate backups: | Not possible to duplicate backups: |
|---|---|
| <ul style="list-style-type: none">♦ from one storage unit to another.♦ from one media density to another.♦ from one server to another.♦ from multiplex to nonmultiplex format.♦ from multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. This is done with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.) | <ul style="list-style-type: none">♦ while the backup is being created (unless you create multiple backup images concurrently during the backup job).♦ while any other backup image is being written to a tape that contains the source primary backup.♦ when the primary backup image is not available.♦ by using the NetBackup scheduler to schedule duplications automatically (unless you use a Vault policy to schedule duplication).♦ of the NetBackup catalogs.♦ when it is a multiplexed image of the following:<ul style="list-style-type: none">– Auspex FastBackup– FlashBackup– NDMP backup– Backups to or from disk type storage units– Nonmultiplexed backups |

If you do multiplexed duplication, be aware of the following:

- ♦ When duplicating multiplexed SQL-BackTrack backups with multiplex mode enabled, it is necessary to duplicate all of the backups in the multiplexed group. This ensures that the fragment order and size is maintained in the duplicate. Otherwise, it is possible that restores from the duplicated backups will not work. A multiplexed group is a set of backups that were multiplexed together during a single multiplexing session.
- ♦ When duplicating multiplexed backups, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups will have a multiplexing factor that is no greater than that used during the original backup.
- ♦ If all backups in a multiplexed group are duplicated to a storage unit that has the same characteristics as the one where the original backup was created, the duplicated group will be identical, with the following exceptions:

- ◆ If end of media (EOM) is encountered on either the source or destination media.
- ◆ If any of the fragments in the source backups are zero length (which can occur if many multiplexed backups start at the same time), during duplication these zero-length fragments are removed.

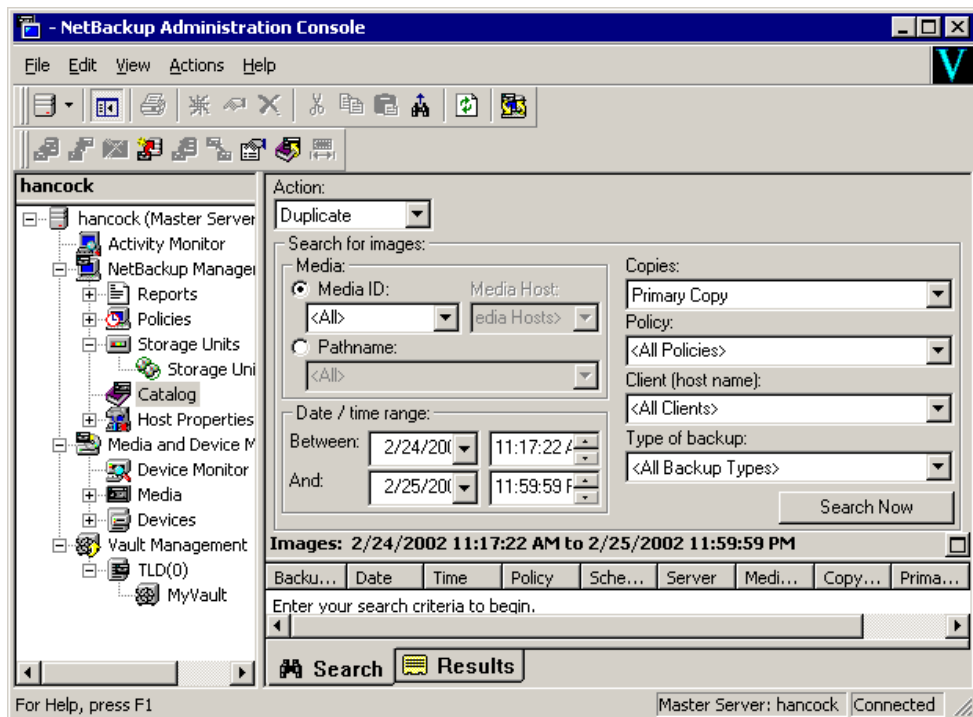
Concurrent Copies through the Catalog Node

Use the following procedure to create concurrent copies of backup images manually through the NetBackup Administration Console Catalog node.

Information similar to the following also is included in the *NetBackup System Administrator's Guide, Volume I*.

▼ To duplicate backup images concurrently through the Catalog node

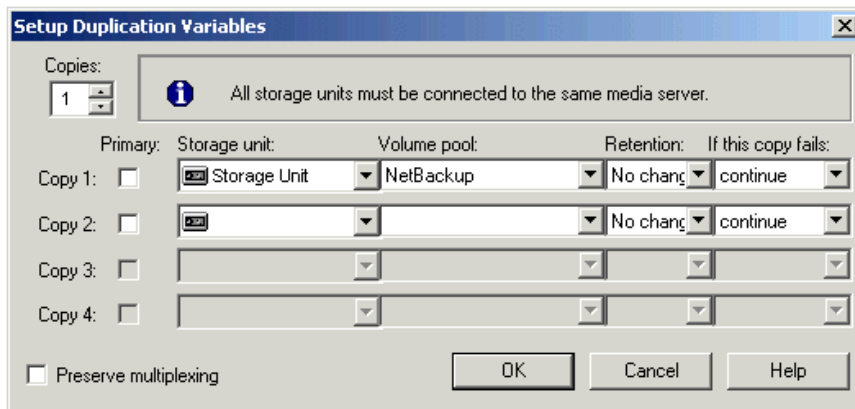
1. In the NetBackup Administration Console, expand **Master Server > NetBackup Management > Catalog**.



2. In the **Action** field, select **Duplicate**.
3. Select the search criteria for the image you want to duplicate, then click **Search Now**.



4. Right-click the image you want to duplicate and select **Duplicate** from the shortcut menu.
5. The Setup Duplication Variables dialog appears.



6. Specify the number of copies to be created.
If enough drives are available, the copies will be created simultaneously. Otherwise, the system may require operator intervention if, for example, four copies are to be created and there are only two drives.
7. If you want one of the duplicated copies to become the primary copy, check the appropriate box.
NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image created during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.
8. Specify the storage unit where each copy will be stored.
If a Media Manager or NDMP storage unit has more than one drive, it can be used for both the source and destination. Network Data Management Protocol (NDMP) storage units are supported only when one copy is created.
9. Specify the volume pool to which each copy will be assigned.
NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as the media ID of the volume that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure a different volume is used.

10. To change the retention level for the copy, select one of the retention level options.

If **No Change** is selected for the retention period, the expiration date is the same for the duplicate and source copies.

If you specify a numeric retention level, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2001 and its retention period is one week, the new copy's expiration date is November 21, 2001.
11. Specify whether the remaining copies should continue or fail if the specified copy fails.
12. If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, select **Preserve Multiplexing**.

If you do not duplicate all the backups in a multiplexed group, the duplicate will have a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

By default, duplication occurs serially using the fewest media mounts and least tape positioning time. Only one backup is processed at a time. If **Preserve Multiplexing** is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated before the multiplexed backups are duplicated.
13. Click **OK** to start duplicating.
14. Click the **Results** tab, then select the duplication job just created to view the job results. For more information, see the *NetBackup System Administrator's Guide, Volume I*.

Concurrent Copies during Basic Duplication

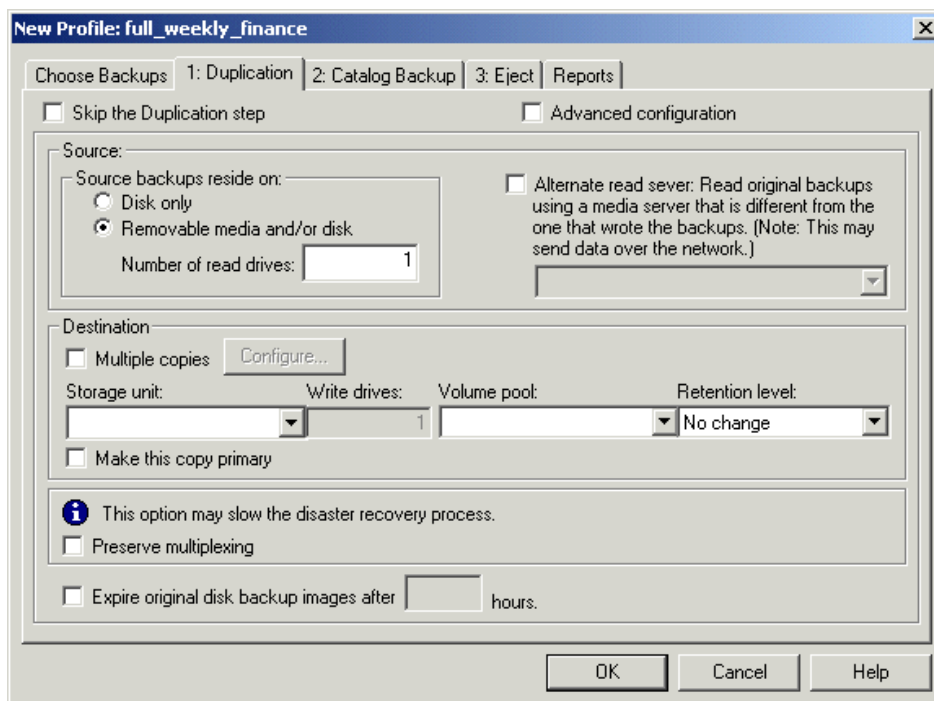
You can create multiple duplicate images concurrently in Vault by selecting either **Multiple Copies** on the basic Duplication tab or **Advanced Configuration** on the basic Duplication tab, which displays the advanced duplication criteria.

You can use the following instructions to create multiple copies concurrently from the basic Duplication tab.

Instructions for configuring duplication in Vault also are included in "[Configuring Duplication](#)" on page 66.



The following is the basic Duplication tab:



▼ To create multiple copies concurrently by using the basic Duplication tab

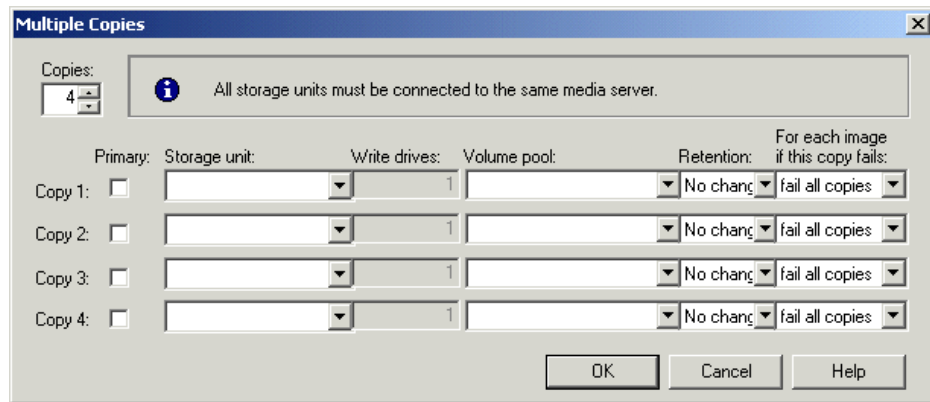
1. Indicate whether the images you want to duplicate reside on disk storage units only or on disk and/or media storage units.
2. Enter the number of drives to be used for reading backup images for duplication.
When you enter a number of read drives, the same number will be entered into the destination Write Drives field. You must have an equivalent number of read and write drives available.
3. To use a media server that is different from the server that wrote the images, check **Alternate Read Server** and select the media server to use.

Note Alternate read servers apply to NetBackup Enterprise Server only.

If robots (or drives) are shared by more than one media server, you can specify a different media server to read the original backups than the media server that wrote the backups.

4. Select **Multiple Copies**, then click **Configure**.

The Multiple Copies dialog appears.



5. Select the number of copies to create.

The number of copies you can choose cannot exceed the number of copies specified in the **Maximum Backup Copies** field for the NetBackup master server. (Configured in **NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes**.) By default, the value is two, which means one original backup and one copy.

6. If you want one of the copies to be the primary copy, select which copy is to be primary.

NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image creating during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.

7. Specify the storage unit to be used for the duplication.

If the Media Manager or NDMP storage unit has more than one drive, the source and destination storage units can be the same. Network Data Management Protocol (NDMP) storage units are supported only when one copy is created.

Note All storage units must be connected to the same media server.

8. Specify a volume pool for each copy.



NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as that of the piece of media that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure a different piece of media is used.

9. Specify the retention level for each copy.

Each image copy can have a separate expiration date. If you do not specify a retention level, it is the same as the primary copy. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify **Use Mappings** for the retention level, the retention period is based on the retention period of backup image copy 1 (for more information, see [“Assigning Multiple Retentions with One Profile”](#) on page 121).

When the retention period expires, information about the expired backup will be deleted from the NetBackup and Media Manager catalog, the volume will be recalled from off-site storage, and the backup image will be unavailable for a restore.

10. Indicate what action is to be taken if a copy fails.

In Vault, if you choose **Fail All Copies**, all copies *of that image* will fail, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault will again try to duplicate the image if the following conditions are true:

- ◆ The image is selected.
- ◆ The Vault profile did not eject the primary backup.

By default, the option is configured to **Fail All Copies** in Vault.

If you choose **Continue** for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted; it is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.

11. Click OK to return to the basic Duplication tab.

12. To preserve multiplexing, select Preserve Multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process speeds up duplication but slows down restores and disaster recovery processes. If the option to preserve multiplexing is selected, the multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session.

Note If the source image is multiplexed and the **Preserve Multiplexing** option is selected, ensure that the destination storage unit configured for each copy has multiplexing enabled. Multiplexing is configured in **NetBackup Management > Storage Units**.

- 13.** Check **Expire Original Disk Backup Images...** and then enter the number of hours after this Vault session completes to expire the disk images.

Use this option to free up space on the disk for subsequent backup images. Be sure you allow enough time for the duplication operation to be completed.

If the duplication of a disk image fails, the disk image will not be expired.

- 14.** After you have completed the dialog, click **OK**.

Concurrent Copies during Advanced Duplication

You can use the following instructions to create multiple copies concurrently from the advanced configuration criteria of the Vault profile **Duplication** tab.

Instructions for configuring duplication in Vault also are included in “[Configuring Duplication](#)” on page 66.

The following shows the **Duplication** tab when **Advanced Configuration** has been selected:

New Profile: full_weekly_finance

Choose Backups | 1: Duplication | 2: Catalog Backup | 3: Eject | Reports

☐ Skip the Duplication step ☒ Advanced configuration

☐ Alternate read server: Read original backups using media servers that are different from the media server that wrote the backups. (Note: This may send data over the network.)

| SOURCE | | DESTINATION | | |
|--------------|-------------|--------------|--------------|-------------|
| Media Server | Read Drives | Storage Unit | Write Drives | Volume Pool |
| | | | | |

New... Delete Change...

i This option may slow the disaster recovery process.

☐ Preserve multiplexing

☐ Expire original disk backup images after hours.

OK Cancel Help



▼ To create multiple copies concurrently by using advanced configuration options

1. On the Duplication tab, select **Advanced Configuration**.

The dialog changes to display more advanced duplication options.

2. To use a server that is different from the server that wrote the images, check **Alternate Read Server**.

Note Alternate read servers apply to NetBackup Enterprise Server only.

If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups.

If you select **Alternate Read Server**, the SOURCE area displays an Alternate Read Server column heading.

3. To add a destination media server and duplication rules for that server, click **New**.

The Duplication Rule dialog appears. If you selected **Alternate Read Server** on the Duplication tab, the Duplication Rule dialog will have fields for both **Source Media Server** and **Alternate Read Server**. If you did not select **Alternate Read Server**, only a **Source Backup Server** field appears.

| Primary | Storage unit | Write drives | Volume pool | Retention | For each image if this copy fails |
|----------------------------------|--------------|--------------|-------------|-----------|-----------------------------------|
| Copy 1: <input type="checkbox"/> | | 1 | | No chang | fail all copies |
| Copy 2: <input type="checkbox"/> | | 1 | | No chang | fail all copies |
| Copy 3: <input type="checkbox"/> | | 1 | | No chang | fail all copies |
| Copy 4: <input type="checkbox"/> | | 1 | | No chang | fail all copies |

4. Select the **Source Media Server**.

5. If you selected **Alternate Read Server** on the Duplication tab, select an **Alternate Read Server**.

Note Alternate read servers apply to NetBackup Enterprise Server only.

The source media server and alternate read server may be the same.

6. Select the number of copies to create.

The number of copies you can choose cannot exceed the number of copies specified in the **Maximum Backup Copies** field for the NetBackup master server. (Configured in **NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes**.) By default, the value is two, which means one original backup and one copy.

7. If you want one of the copies to be the primary copy, select which copy is to be primary.

NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image creating during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured continue as the fail option, the first successful copy is the primary copy.

8. Specify the storage unit to be used for the duplication.

If the Media Manager or NDMP storage unit has more than one drive, the source and destination storage units can be the same. Network Data Management Protocol (NDMP) storage units are supported only when one copy is created.

Note All storage units must be connected to the same media server.

9. Specify a volume pool for each copy.

NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as that of the piece of media that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure a different piece of media is used.

10. Specify the retention level for each copy.

Each image copy can have a separate expiration date. If you do not specify a retention level, it is the same as the primary copy. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify **Use Mappings** for the retention level, the retention period is based on the retention period of backup image copy 1 (for more information, see [“Assigning Multiple Retentions with One Profile”](#) on page 121).



When the retention period expires, information about the expired backup will be deleted from the NetBackup and Media Manager catalog, the volume will be recalled from off-site storage, and the backup image will be unavailable for a restore.

11. Indicate what action is to be taken if a copy fails.

In Vault, if you choose **Fail All Copies**, all copies *of that image* will fail, independent of the success or failure of other image copy operations. The next time the vault profile runs, Vault will again try to duplicate the image if the following conditions are true:

- ◆ The image is selected.
- ◆ The Vault profile did not eject the primary backup.

By default, the option is configured to **Fail All Copies** in Vault.

If you choose **Continue** for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted; it is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.

12. Indicate whether you want to preserve multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process speeds up duplication but slows down restores and disaster recovery processes. If the option to preserve multiplexing is selected, the multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session.

Note If the source image is multiplexed and the **Preserve Multiplexing** option is selected, ensure that the destination storage unit configured for each copy has multiplexing enabled. Multiplexing is configured in **NetBackup Management > Storage Units**.

13. Check **Expire Original Disk Backup Images... and then enter the number of hours after this Vault session completes to expire the disk images.**

Use this option to free up space on the disk for subsequent backup images. Be sure you allow enough time for the duplication operation to be completed.

If the duplication of a disk image fails, the disk image will not be expired.

14. Click **OK.**

The reports for each profile are configured in the profile dialog **Reports** tab. When you configure a Vault profile, you specify which reports should be generated, when they should be generated, and how and to whom they should be distributed.

After reports are generated and distributed, you can view and print them until the Vault logs for that session are deleted.

For more information, see the following:

- ◆ [“Generating Reports”](#) on page 159
- ◆ [“Consolidating Reports”](#) on page 162
- ◆ [“Viewing Reports”](#) on page 163
- ◆ [“Vault Report Types”](#) on page 164

Related Topics

- ◆ [“Configuring Reports”](#) on page 91
- ◆ [“Vault Session Log Files”](#) on page 184

Generating Reports

If the reports for a profile are configured as immediate, the reports are generated when the profile runs. If the reports for a profile are deferred, you can use one of the following methods to generate the reports after the profile runs:

- ◆ By invoking the **Deferred Eject** option in the Administration Console and then selecting **Generate Reports After Eject**.
- ◆ By using the Vault Operator Menu interface.
- ◆ By using the `vlteject` command.
- ◆ By using a Vault policy that invokes the `vlteject` command.

When you generate reports, the reports that are selected on the profile dialog **Reports** tab are generated and distributed to the destinations specified.



Reports can be generated for one session or for multiple sessions. Generating reports and ejecting media from more than one vault session is known as *consolidating* your reports and ejections. For example, you may duplicate images daily but eject media and generate reports only at the end of the week.

Related Topics

- ◆ [“Ejecting Media by Using the NetBackup Administration Console”](#) on page 107
- ◆ [“Configuring Reports”](#) on page 91
- ◆ [“Reports that Depend on Eject”](#) on page 96

Generating Reports by Using the Vault Operator Menu

You can use the Vault Operator Menu to generate reports.

▼ To generate reports by using the Vault Operator Menu

1. Start the Vault Operator Menu by invoking the `vltopmenu` command.
2. If necessary, select a profile.
3. Select one of the following options:
 - ◆ **Run Reports for This Session**
 - ◆ **Run Individual Reports**
 - ◆ **Consolidate All Reports**
 - ◆ **Consolidate All Reports and Ejects**

Consolidating reports and ejects also ejects media.
4. Continue as prompted by the Vault Operator Menu.

Generating Reports by Using the `vlteject` Command

You can use the `vlteject` command with the `-report` option to generate reports from the command line. The following is the syntax for the command that generates all reports that have not yet been generated:

```
vlteject -report
```

You also can specify a robot, vault, profile, or session for which to generate reports.

If the corresponding eject process has been completed, pending reports from the sessions selected are generated and distributed. The reports will not be generated again if `vlteject` is run again.

If eject has not been completed, the subset of reports that do not depend on completion of eject will be generated. These reports will be generated again if `vlteject` is run again.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name] [-robot
robot_name] [-vault vault_name [-sessionid id]] [-auto y|n]
[-eject_delay seconds]
```

The `vlteject` command resides in the following directory:

UNIX: `/usr/opensv/netbackup/bin`

Windows: `install_path\NetBackup\bin`

For more information about the `vlteject` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To generate reports by using the `vlteject` command

1. In a terminal or command window, change to the directory in which the `vlteject` command resides.
2. Invoke the command, using the appropriate options and parameters.

Generating Reports by Using a Vault Policy

You can use a Vault policy to generate reports for Vault sessions that have been completed already and for which reports have not been generated. In the Vault policy, specify Vault as the policy type, do not specify clients, and specify the `vlteject` command with the `-report` option on the **Backup Selections** tab.

You also can specify a robot, vault, profile, or session for which to generate reports.

If the corresponding eject process has been completed, pending reports from the sessions selected are generated and distributed. The reports will not be generated again if `vlteject` is run again.

If eject has not been completed, the subset of reports that do not depend on completion of eject will be generated and distributed. These reports will be generated again if `vlteject` is run again.

The following is the format of the `vlteject` command:



```
vlteject -eject -report -preview [-profile profile_name] [-robot  
robot_name] [-vault vault_name [-sessionid id]] [-auto y|n]  
[-eject_delay seconds]
```

The `vlteject` command resides in the following directory:

UNIX: `/usr/opensv/netbackup/bin`
Windows: `install_path\NetBackup\bin`

For more information about the `vlteject` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* manual.

▼ To create a Vault policy that generates reports

1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Click the **New Policy** button.
The Add a New Policy dialog appears.
3. Enter a unique name for the new policy in the Add a New Policy dialog.
4. Click **OK**.
The Add New Policy dialog appears.
5. On the **Attributes** tab, select **Vault** as the policy type.
6. On the **Schedules** tab, click **New** to create a new schedule.
The type of backup defaults to **Automatic Vault**.
7. Complete the schedule.
8. Bypass the **Client** tab (clients are not specified for Vault jobs).
9. On the **Backup Selections** tab, enter the `vlteject` command with the `-report` option and any other appropriate options.
10. Click **OK**.

Consolidating Reports

Note If you consolidate reports, you should also consolidate ejects.



You can generate reports and eject media from more than one vault session, which is known as *consolidating* your reports and ejections. For example, you may duplicate images daily but eject media and generate reports only at the end of the week. To do so, specify deferred reports on the **Reports** tab and deferred eject on the **Eject** tab for each profile for which you want to consolidate reports. Then, eject the media and generate the reports using one of the methods in “[Ejecting Media](#)” on page 106.

When you generate the reports, you select the robot, vault, or profile sessions for which reports were deferred (that is, for reports that are pending).

Each report from all participating sessions will be concatenated into a single report. The reports are then distributed to all destinations specified on the **Reports** tab of all participating profiles. For example, a consolidated Picking List for Robot will be created, and it will include the Picking List for Robot from all sessions whose profile has Picking List for Robot selected on the **Reports** tab.

If eject has not been completed, the subset of reports that do not depend on completion of eject will be generated; these reports will be generated again if deferred reports are run again.

If you consolidate reports and also rename reports, use the same customized report title for all profiles whose reports will be consolidated. The customized report title is printed on the report and appears in the e-mail subject line if you e-mail the reports.

Related Topics

- ◆ “[Consolidating Ejects](#)” on page 111
- ◆ “[Reports that Depend on Eject](#)” on page 96
- ◆ “[Configuring Eject](#)” on page 85
- ◆ “[Configuring Reports](#)” on page 91

Viewing Reports

You can use the NetBackup Administration Console to view and print reports for Vault sessions for which reports have already been generated. You can only view reports if the session directory for that vault still exists. Only some of the reports will be valid; for example, the picking list reports are only valid on the date they were generated.

▼ To view Vault reports

1. Select **NetBackup Management > Reports > Vault Reports**.
2. Select one of the Vault reports or report types.



When you select a report or report type, the Reports window is displayed. The Reports window includes a Report Settings area and a report contents window.

3. Enter or select the appropriate values for the report you want to generate.
Usually, you must specify a profile and a session ID; you also may have to specify a date range or time period.
4. Click **Run Report**.
5. To print the report, click **File > Print**.

Vault Report Types

The following reports and types of reports are available in Vault:

- ◆ [Reports for Media Going Off-Site](#)
- ◆ [Reports for Media Coming On-Site](#)
- ◆ [Inventory Reports](#)
- ◆ [Container Inventory Report](#)
- ◆ [Recovery Report for Vault](#)
- ◆ [Lost Media Report](#)
- ◆ [Non-vaulted Images Exception Report](#)
- ◆ [Iron Mountain FTP File](#)

Reports for Media Going Off-Site

The reports for media going off-site show the volumes that have been ejected from the robot and will be transported off-site. They vary in the amount of detail included in each report.

Picking List for Robot

The Picking List for Robot report shows the volumes ejected from the robot that should be transported off-site. This report is sorted by media ID and should be used by the operations staff as a checklist for media that has been ejected from the robots. You can save the report for tracking purposes, or you can reprint it as long as the session directory still exists.

The information in the report depends on whether the vault uses containers or slots.

Picking List for Robot Report Headings

| Column | Description |
|--------------|--|
| #IMAGES | The number of images on the volume. For Vault catalog backup volumes, displays zero. |
| CONTAINER ID | The ID of the container in which the volume will reside in the vault. |
| EXPIRATION | Date when the images on the volume expire. For Vault catalog backup volumes, displays the date calculated as a return date during the volume assignment. |
| KBYTES | The size in kilobytes of images on the volume. For Vault catalog backup volumes, displays the notation NB Catalog. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| SLOT ID | The ID of the slot in which the volume will reside at the off-site vault. |

Distribution List for Vault

The Distribution List for Vault report shows the volumes that have been ejected from the robot and will be transported off-site. This report is sorted by off-site slot number and should accompany the media that is destined for the off-site vault. The vault vendor should use this report to verify that all the volumes listed were actually received.

The information in the report depends on whether the vault uses containers or slots.

Distribution List for Vault Report Headings

| Column | Description |
|--------------|--|
| #IMAGES | The number of images on the volume. For Vault catalog backup volumes, displays zero. |
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| EXPIRATION | Date when the images on the volume expire. For Vault catalog backup volumes, displays the date calculated as a return date during the volume assignment. |



Distribution List for Vault Report Headings (continued)

| Column | Description |
|-------------|--|
| KBYTES | The size in kilobytes of images on the volume. For Vault catalog backup volumes, displays the notation NB Catalog. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| RETURN DATE | The date the container should be returned from the off-site vault. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Detailed Distribution List for Vault

This report is similar to the Picking List for Robot and Distribution List for Vault reports except that detailed information is listed for each media. Because backup jobs can span volumes, it is possible that detailed listings of a backup job appear on more than one volume. This report is useful at a disaster recovery site. VERITAS recommends that you send this report off-site.

The information in the report depends on whether the vault uses containers or slots.

Detailed Distribution List for Vault Report Headings

| Column | Description |
|--------------|--|
| #IMAGES | The number of images on the volume. For Vault catalog backup volumes, displays zero. |
| BACKUP ID | Identifier that NetBackup assigns when it performs the backup. |
| BACKUP TIME | When the backup occurred. |
| CLIENT | Name of the client that was backed up. |
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| EXPIRATION | Date when the images on the volume expire. For Vault catalog backup volumes, displays the date calculated as a return date during the volume assignment. |
| KBYTES | The size in kilobytes of images on the volume or the size in kilobytes of the image fragments on the volume. TIR indicates a true image restore image. For Vault catalog backup volumes, displays the notation NB Catalog. |

Detailed Distribution List for Vault Report Headings (continued)

| Column | Description |
|----------|---|
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| PARTIAL? | Partial images on the volume. The field displays: <ul style="list-style-type: none"> ♦ Complete if the image is not partial (that is, does not span volumes). ♦ PARTIAL(<i>x</i>) if it is a partial images (<i>x</i> is the fragment number). ♦ EXTRA if the images does not belong to the session. |
| POLICY | Name of the policy that was used to back up the client. |
| SCHEDULE | Name of the schedule that was used to back up the client. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Summary Distribution List for Vault

This report is similar to the Detailed Distribution List for Vault report, except that the entry for each piece of media will list only a unique client, policy, schedule and date. That is, if multiple backup jobs for a given client, policy and schedule (usually seen with RDBMS backups or SAP backups) are written to the same volume on the same date, only one line of information will be printed out on this report. The Detailed Distribution List would show each of these backup jobs as a separate entry, which may generate a very long report. The Summary Distribution List for Vault report summarizes the information and presents it in a more compact form. This report is also very useful for disaster recovery situations; we recommend that you send this report off-site.

The information in the report depends on whether the vault uses containers or slots.

Summary Distribution List for Vault Report Headings

| Column | Description |
|--------------|--|
| #IMAGES | The number of images on the volume. For Vault catalog backup volumes, displays zero. |
| BACKUP TIME | When the backup occurred. |
| CLIENT | Name of the client that was backed up. |
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |



Summary Distribution List for Vault Report Headings (continued)

| Column | Description |
|------------|--|
| EXPIRATION | Date when the images on the volume expire. For Vault catalog backup volumes, displays the date calculated as a return date during the volume assignment. |
| KBYTES | The size in kilobytes of images on the volume. For Vault catalog backup volumes, displays the notation NB Catalog. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| POLICY | Name of the policy that was used to back up the client. |
| SCHEDULE | Name of the schedule that was used to back up the client. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Reports for Media Coming On-Site

The reports for media coming on-site show the volumes that are being requested back from the off-site vault. Vault will not generate these reports until the media have been ejected for the current Vault session.

Picking List for Vault

The Picking List for Vault report shows the volumes that are being requested back from the off-site vault. This report should be sent off-site to the vault vendor. Volumes are listed on this report because Vault determined that they are in an off-site volume group and that all images have expired. When Vault identifies these volumes, it changes the Date Requested field within the Media Manager description field for the media. It then prints out the media ID on this report along with the date requested.

Expired media only appear on the report generated on the date the media expire or the next time the report is generated. If for some reason expired media does not appear on the Picking List for Vault, they will be listed on the Lost Media report.

A slot at the off-site vault from which an expired volume is recalled will be available for use one day after the volume has been physically returned to the robot.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

The information in the report depends on whether the vault uses containers or slots.

Picking List for Vault Report Headings

| Column | Description |
|--------------|---|
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| DENSITY | Density of the volume. |
| LAST MOUNT | The date the volume was last mounted. For Vault catalog backup volumes, displays the notation NB Catalog. |
| LAST SID | The last session ID of the profile that accessed this volume. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| REQUESTED | Date when the volume is requested back from the off-site vault. |
| RETURN DATE | The date the container should be returned from the off-site vault. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Distribution List for Robot

The Distribution List for Robot report shows the volumes that are being requested back from the off-site vault. This report is identical to the Picking List for Vault, except that it has a different report title. Retain this report on-site to use as a checklist for the media returned from the off-site vault.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

The information in the report depends on whether the vault uses containers or slots.

Distribution List for Robot Report Headings

| Column | Description |
|--------------|---|
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| DENSITY | Density of the volume. |



Distribution List for Robot Report Headings (continued)

| Column | Description |
|-------------|---|
| LAST MOUNT | The date the volume was last mounted. For Vault catalog backup volumes, displays the notation NB Catalog. |
| LAST SID | The last session ID of the profile that accessed this volume. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| REQUESTED | Date when the volume is requested back from the off-site vault. |
| RETURN DATE | The date the container should be returned from the off-site vault. |
| ROBOT | The robot from which the volumes were ejected. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Inventory Reports

The inventory reports show the location of the media. These reports are not generated until the media have been ejected.

If you use the NetBackup Administration Console to display an inventory report, you must select a profile that ejects media. Also, select the most recent session for that profile so the most recent data is reported.

Vault Inventory

The Vault Inventory (or Inventory List for Vault) report shows all media that are off-site at the vault vendor and media being sent off-site. This list is generated by checking the description field for the media, the volume pool, and the off-site volume group. VERITAS recommends that you send this report to your vault vendor so they can verify that they have the volumes that Vault indicates are at the vault vendor.

The information in the report depends on whether the vault uses containers or slots.

Vault Inventory Report Headings

| Column | Description |
|----------|---|
| ASSIGNED | The date when the volume was assigned by NetBackup Media Manager. For Vault catalog backup volumes, displays the notation NB Catalog. |



Vault Inventory Report Headings (continued)

| Column | Description |
|--------------|---|
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| EXPIRATION | Date when the images on the volume expire. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Off-site Inventory

The Off-site Inventory (or Full Inventory List for Vault) report includes the information in the Vault Inventory report and also includes any volumes that have been requested back from the off-site vault vendor (that is, volumes in transit). Usually, this report is not generated on a daily basis. Rather, the Inventory List for Vault report is sent to the vault vendor to perform verification.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

The information in the report depends on whether the vault uses containers or slots.

Off-site Inventory Report Headings

| Column | Description |
|--------------|---|
| ASSIGNED | The date when the volume was assigned by NetBackup Media Manager. For Vault catalog backup volumes, displays the notation NB Catalog. |
| CONTAINER ID | The ID of the container in which the volume resides in the vault. (Container vaulting only.) |
| EXPIRATION | Date when the images on the volume expire. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| REQUESTED | The date the volume was requested to be returned from the off-site vault. |
| SLOT ID | The ID of the slot in which the volume resides in the vault. (Slot vaulting only.) |



All Media Inventory

The All Media Inventory (or Complete Inventory List for Vault) report shows all volumes in the off-site volume pool.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

Note Volumes within the off-site volume pool must belong to either the off-site volume group or the robotic volume group or they will not appear on this report.

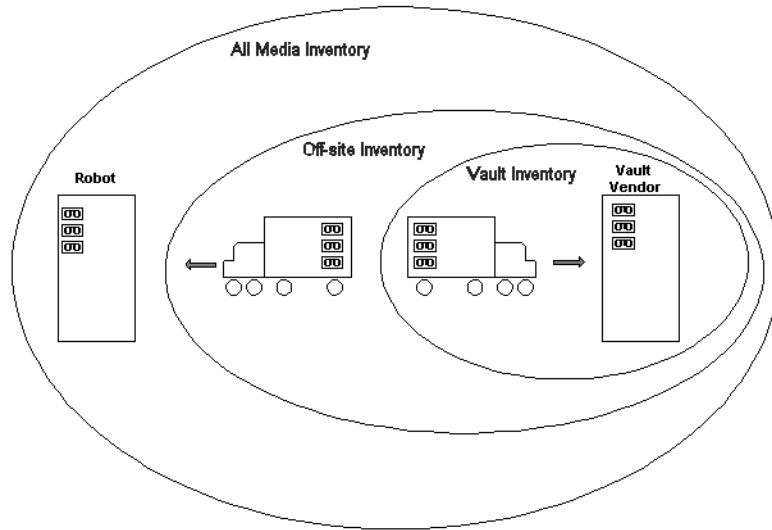
The information in the report depends on whether the vault uses containers or slots.

All Media Inventory Report Headings

| Column | Description |
|--------------|---|
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| EXPIRATION | Date when the images on the volume expire. |
| LOCATION | Where the volume resides, the robot (R) or vault (V). |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| REQUESTED | The date the volume was requested to be returned from the off-site vault. |
| SID | The ID of the session that duplicated and/or ejected this volume. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |

Graphical Representation of Inventory Reports Scope

The following illustration shows the different scopes of the reports:



Container Inventory Report

The Container Inventory Report shows all the containers configured in your vaulting environment, the return date of each container, and the media that are in each container. Alternatively, you can specify a container ID to generate a report of the media in a specific container.

Generate this report only if you vault your media in containers. Reports will not show container information until after you add container and media IDs in Vault. Media are removed logically from a container when they are injected back into the robot.

Container Inventory Report Headings

| Column | Description |
|--------------|---|
| CONTAINER ID | The ID of the container in which the volume resides in the vault. |
| LAST SID | The last session ID of the profile that accessed this volume. |
| MEDIA ID | ID of the media ID that are in the container. |



Container Inventory Report Headings (continued)

| Column | Description |
|-------------|--|
| REQUESTED | Date when the container is requested back from the off-site vault. |
| RETURN DATE | The date the container should be returned from the off-site vault. |
| ROBOT | The robot from which the volumes were ejected. |

Recovery Report for Vault

The Recovery Report for Vault shows all policies defined on a NetBackup master server and all media that is required for restores between a given set of dates. The report displays the date range to which the images on the media apply.

This report includes the three most recent NetBackup catalog volumes that are currently off-site. For the NetBackup catalog media to be listed in this section, their volume group must match the volume group specified in the off-site volume group. Only NetBackup catalog media that are assigned will appear on this report.

Sending the Recovery Report to the vault vendor on a regular basis will help with disaster recovery efforts. If the master server is destroyed by a disaster, you will not be able to generate a Recovery Report to determine which volumes to request from the vault vendor. Therefore, it is very important that the vault vendor have a copy of the Recovery Report.

The information in the report depends on whether the vault uses containers or slots.

Recovery Report for Vault Fields

| Field | Description |
|---------------|--|
| BACKUP POLICY | Name of the policy that was used to back up the client. |
| CLIENT | Name of the client that was backed up. |
| DATE | Date when the original backup occurred. |
| EXPIRATION | Date when the catalog backup expire. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| SCHEDULE | Name of the schedule that was used to back up the client. |
| WRITTEN | The date and time the catalog backup was written to this volume. |

Lost Media Report

The Lost Media report lists expired media that has not been returned from the off-site vault vendor. Media can get stranded at the off-site vault for various reasons, as follows:

- ◆ Frozen backup media never expires. Media that does not expire will not appear on the Picking List for Vault and will not be recalled from the vault.
- ◆ A volume appears on the Picking List for Vault only once. If a volume from that report is missed and is not returned to the robot, it will never again be listed for recall.

You must generate the Lost Media Report; it is not generated when you eject media. You do not have to configure your profiles for the Lost Media Report. Usually, media included in the Lost Media Report should be returned from off-site and injected back into the appropriate vault in the robot.

A good practice is to run the Lost Media Report periodically, such as weekly or monthly (depending on your operations). The Lost Media Report will list media that expired and should have been returned on-site and injected back into a robot for reuse.

Lost Media Report Headings

| Column | Description |
|--------------|--|
| DENSITY | Density of the volume. |
| LAST MOUNT | The date the volume was last mounted. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| REQUESTED | Date when the volume is requested back from the off-site vault. |
| VAULT | The vault to which the volume belongs. |
| VOLUME GROUP | The volume group to which the volume is assigned. |

Non-vaulted Images Exception Report

The Non-vaulted Images report lists images that are not in an off-site volume pool (that is, images that were not duplicated) and media that were not ejected and therefore were not transferred to the off-site vault vendor. When generated as part of a scheduled Vault job, the Non-vaulted Images report uses the same time window as the Vault session and includes information for that specific profile; if generated manually, you can specify a date range by one of the following methods:



- ◆ Specifying a calendar date
- ◆ Specifying a range of days beginning x days ago
- ◆ Specifying a session date range beginning with the start time of the session and using the profile's time window

Non-vaulted Images Report Headings

| Column | Description |
|--------------|---|
| ASSIGNED | The date when the volume was assigned by NetBackup Media Manager. |
| BACKUP ID | Identifier that NetBackup assigns when it performs the backup. |
| CLIENT | Name of the client that was backed up. |
| CREATED | The date the volume was created (original backup or duplicated). |
| EXPIRATION | Date when the images on the volume expire. |
| MEDIA ID | Media ID that is assigned when the volume is added to Media Manager. |
| POLICY | Name of the policy that was used to back up the client. |
| SCHEDULE | Name of the schedule that was used to back up the client. |
| SLOT ID | The ID of the slot in which the volume resides in the off-site vault. |
| VOLUME GROUP | The volume group to which the volume is assigned. |
| VOLUME POOL | The volume pool to which the volume is assigned. |

Iron Mountain FTP File

If Iron Mountain is your vault vendor, you can configure Vault to produce an Iron Mountain Electronic Format report, which is a file that can include the following reports:

- ◆ Picking List for Vault
- ◆ Distribution List for Vault
- ◆ Off-site Inventory Report (if you are vaulting in slots)
- ◆ Container Inventory Report (if you are vaulting containers)
- ◆ Recovery Report



The reports included in the file depend on your selections on the **Reports** tab of the profile dialog; you must select a report so that it will appear in the Iron Mountain report file.

The report will be in a format that Iron Mountain's automated vaulting mechanism can read and contain the information they require. You can use the file transfer protocol (FTP) to send the report file to Iron Mountain electronically, and they use it to update their vaulting mechanism automatically.

Before you send the report to Iron Mountain, you should verify that the volumes ejected match the Distribution List for Vault. You should contact Iron Mountain to determine where and when to send the report.



Administering Vault

The following sections provide information about performing the tasks of managing your Vault configuration.

- ◆ [“Setting Up E-Mail”](#) on page 179
- ◆ [“Administering Access to Vault”](#) on page 180
- ◆ [“Printing Vault and Profile Information”](#) on page 182
- ◆ [“Copying a Profile”](#) on page 182
- ◆ [“Moving a Vault to a Different Robot”](#) on page 183
- ◆ [“Changing Volume Pools and Groups”](#) on page 184
- ◆ [“Vault Session Log Files”](#) on page 184
- ◆ [“General Operational Issues”](#) on page 187

Setting Up E-Mail

Depending on your computing environment, you may have to configure NetBackup or your computing environment so that notification e-mail from NetBackup functions properly.

On UNIX systems, the `mail` or `mailx` command is used to send e-mail. If `mail` or `mailx` is installed, NetBackup uses it to send email. If not installed, you must install one of those mail services and configure your environment so it functions correctly.

On Windows systems, it may be necessary to configure NetBackup for e-mail by using the `nbmail.cmd` script (in `install_path\VERITAS\NetBackup\bin`) on the system on which the NetBackup master server is installed. For e-mail notifications, NetBackup passes the e-mail address, subject, and message to the script. The script then uses the mailing program specified in the script to send e-mail.

For instructions on configuring the script to use a third-party e-mail client, see the comments in the `nbmail.cmd` script. Default NetBackup behavior: `nbmail.cmd` does not send e-mail.



Note If you use the Blat e-mail client to deliver e-mail on Windows systems, include the `-mime` option on the `blat` command in the `nbmail.cmd` script so that the Vault reports are mailed correctly.

For more information about configuring e-mail for NetBackup, see the *NetBackup System Administrator's Guide, Volume I*.

Administering Access to Vault

NetBackup provides two mutually exclusive methods for controlling user access:

- ◆ **Access Management.** Access Management lets you control access to NetBackup by defining user groups and granting explicit permissions to these groups. Configuring user groups and assigning permissions is done using Access Management in the NetBackup Administration Console. Access Management is the newest and will be the preferred method in future NetBackup releases.
- ◆ **Enhanced Authorization and Authentication.** Enhanced authentication allows each side of a NetBackup connection to verify the host and user on the other side of the connection. By default, NetBackup runs without enhanced authentication. Enhanced authorization determines if authenticated users (or groups of users) have NetBackup administrative privileges. By default, NetBackup gives administrative privileges to UNIX root administrators or Windows system administrators on NetBackup servers.

If both Access Management and Enhanced Authorization and Authentication are configured, Access Management takes precedence.

For information about configuring and using these methods to control access to Vault, see the *NetBackup System Administrator's Guide, Volume II*.

Vault Operator User Group

NetBackup Access Management is used to define user groups, specify which actions each user group can perform, and assign users to those user groups. Each user group can perform only the actions explicitly granted and no others.

When Vault is installed and licensed, the NetBackup 5.0 release includes a Vault Operator user group that has permission to perform the operator actions necessary for the Vault process. In NetBackup Access Management terminology, the Vault Operator user group has the following permissions:

Vault Operator Permission Sets Defaults

| Permission Sets | Permissions | Vault Operator |
|-----------------|-----------------|----------------|
| Operate media | Browse media | X |
| | Read media | X |
| | Inject media | X |
| | Eject media | X |
| | Move media | X |
| | Assign media | X |
| | Deassign media | X |
| | Update database | X |
| Read report | Browse report | X |
| | Read report | X |
| Operate robot | Browse robot | X |
| | Read robot | X |
| | Inventory robot | X |

These permissions are granted only in the scope of actions performed in Vault. For example, the Vault Operator group has permission to update databases — but only to the extent allowed by Vault, such as ejecting media which changes volume group information for the volume ejected. As defined in the default permission sets, the Vault Operator cannot use the NetBackup Administration Console to change database information that is not related to the operate media actions.

If you use Access Management to administer access by using the default Vault Operator group, those permission sets and permissions apply regardless of whether the actions are invoked from the Vault Operator Menu or the NetBackup Administration Console.



A NetBackup Security Administrator (a user group defined within NetBackup Access Management) can use Access Management to add users to the Vault Operator group and change the permission sets and permissions of the Vault Operator group. A Security Administrator also can create new user groups to define new roles.

Because you can change which actions user groups can perform, the Vault documentation cannot specify which actions are or are not allowed by Access Management. If an action cannot be performed because of access management restrictions, NetBackup Administration Console messages will explain the restriction.

For information about installing Access Management components and using Access Management, see “Access Management” in the *NetBackup System Administrator’s Guide, Volume II*.

Caution Giving operators access to the Vault Operator Menu also gives operators the capability to change report destinations. If you do not want your operators to view reports and change report destinations, do not give them access to the Vault Operator Menu. For example, you may not want your operators to see the Recovery Report or to be able to change to whom reports are e-mailed.

Printing Vault and Profile Information

You can print a list of the information (robots, vaults, or profiles) that is currently displayed in the Administration Console Details Pane. From the **File** menu, choose **Print** or click the **Print** icon on the toolbar.

Copying a Profile

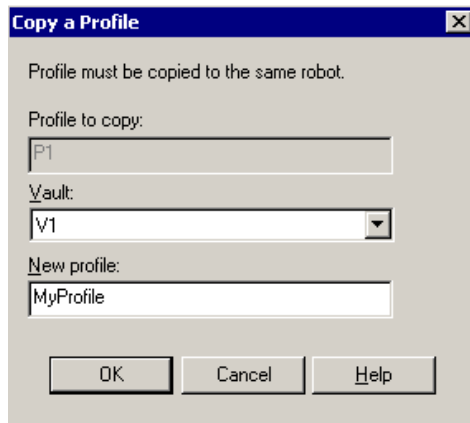
If you want to create a profile that is similar to another profile, you can copy the existing profile, rename it, and then change the attributes.

Note The new profile must belong to the same robot as the original profile.

▼ To copy a profile

1. Highlight the profile you want to copy.
2. Open the **Actions** menu and select **Copy Profile**.

The Copy a Profile dialog displays.



3. From the dropdown list under the **Vault** field, choose the vault in which to place the new profile.
4. Enter a new name for the profile.
5. Click **OK**.

Moving a Vault to a Different Robot

A vault is associated with (that is, belongs to) a specific robot. However, you can change the robot to which the vault belongs. To do so, right-click on the robot and select **Change**. Then complete the dialog, specifying another robot for the vault, and click **OK**.

Note All vaults that had been associated with the previous robot will now be associated with the new robot chosen in the dialog. Some profile configurations may be invalid under the new robot; for example, if the previous robot was associated with a media server that the new robot is not associated with, the configuration will be invalid.

Robots are configured in NetBackup through Media Manager. The action described here does not change the configuration of a robot in Media Manager.



Changing Volume Pools and Groups

Caution If you have media in your off-site vault, VERITAS recommends that you do not change or rename your off-site volume group(s) or off-site volume pool(s). If you begin using new volume pools and groups in your Vault profiles, the reports that recall expired media will not include the old groups and pools.

If you change to a new off-site volume group and/or off-site volume pool(s), you can ensure that media are recalled by configuring a profile that only generates the reports needed to recall media, as follows:

- ◆ Configure a vault that uses the old off-site volume group
- ◆ Configure a profile in that vault that does the following:
 - ◆ Selects no images (that is, configure the Choose Backups step so that no backup images are selected)
 - ◆ Skips the Duplication and Catalog Backup steps
 - ◆ Specifies the old volume pools in the volume pool list on the Eject step
 - ◆ Generates only the Picking List for Vault and Distribution List for Robot reports
- ◆ Schedule the profile to run on a regular basis

Media in the old off-site volume group and in the old volume pool(s) will be recalled from off-site storage as they expire. After each volume is recalled and injected back into the robot, change its volume pool and group to the new ones (if a volume is returned to a scratch volume pool from which all media are allocated, you do not have to change the volume pool).

After all media in those volume pools and groups are recalled, you can delete the vault, volume group, and volume pools.

Vault Session Log Files

Vault generates session logs and debug logs. The session logs can help you keep track of Vault processes.

Related Topics

- ◆ [“Debug Logs”](#) on page 202

Session Logs

The session directory generated for each vault session collects information for the session in two log files. The `detail.log` file contains a step-by-step account of each action performed for the session. Some of the information in `detail.log` is also recorded in the NetBackup log files. The `summary.log` file contains a brief description of the vault session, and the results of the session. If e-mail notification is enabled, the information in this file is appended to the e-mail.

The `detail.log` has information about the number of images selected by a particular session. In addition, it should record information (during the duplication step) about the total number of images and the number of images duplicated. If these numbers do not match, it means that some images were not duplicated. The log should contain information about which images were not duplicated, either because they were duplicated in a prior session or because the duplication failed for some reason. The actual images selected by the session will show up only if a higher debug level (level 5) is used.

Vault maintains its session log files for a particular session in the directory for that session. The directory is located in the following path:

UNIX: `/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx`

Windows: `install_path\NetBackup\vault\sessions\vault_name\sidxxx`

vault_name is the name of the vault used for the session and *xxx* is the unique session identifier that Vault assigns to each vault session. The session ID starts at 1 the first time Vault runs and is incremented by 1 for each new session. The session identifier for a Vault session can be found by looking at the Activity Monitor entry for that session.

The following table describes the session logs.

Vault Session Log Files

| Name | Purpose |
|-------------------------------|--|
| <code>duplicate.log.nn</code> | Progress information for duplication operations; generated by the <code>-L</code> option of the <code>bpduplicate</code> command. |
| <code>preview.list</code> | Summary of images to be duplicated if Duplication step is configured, or ejected, if Eject step is configured and Duplication step is not. |
| <code>image.list</code> | Lists all images and partial images for a session. |
| <code>detail.log</code> | Summary of each action performed for a Vault session. |



Vault Session Log Files (continued)

| Name | Purpose |
|---------------------------------------|---|
| <code>summary.log</code> | Brief description of the Vault session and its results. If e-mail notification is enabled, data in this log file is appended. |
| <code>vltrun.output_vault_name</code> | Output file that shows the progress of the session. |

Setting the Duration of Vault Session Files

Vault's session files are stored in the following directory:

UNIX: `/usr/opensv/netbackup/vault/sessions`

Windows: `install_path\NetBackup\vault\sessions`

You can configure the length of time NetBackup retains these files by using the NetBackup Administration Console.

▼ To set the length of time to keep working files

1. In the NetBackup Administration Console, select **Host Properties**.
2. Select **Master Server** under **Host Properties**.
3. In the right pane, right-click the master server and choose **Properties**.
4. Select **Global Attributes**.
5. In the **Delete vault logs** field, set the length of time after which to delete the Vault working files.

When the set time has elapsed, the entire `sidxxx` directory is deleted.

You should plan to retain each `sidxxx` directory at least as long as the period of time over which you plan to span consolidated ejects. We suggest that you keep these directories at least a week longer than the consolidation span. If the `sidxxx` directory has been deleted, Vault will be unable to eject tapes or generate reports from that session.

Output from the Vault Driver

The vault driver, `vltrun`, produces an output file that shows the progress of the vault session. The output file (`vltrun.output_vault_name`) resides in each vault session directory.



As each step of a vault session completes, the result of the step is written to this file. If the session fails, the file only contains information up through the last step successfully completed.

General Operational Issues

General operational issues describe issues that you should be aware of when configuring and using Vault.

Vaulting Storage Migrator Files

Files migrated by VERITAS Storage Migrator are moved to secondary storage, and a pointer to each file remains in the filesystem. Consequently, a regular backup of the filesystem will save only the file pointer information. If Vault duplicates those files, only the file pointer information is vaulted. To ensure that the actual file is vaulted, you should also vault a copy of the Storage Migrator media.

You can use the `vlt_ejectlist_notify` script to add the Storage Migrator media to the list of tapes to be vaulted. For more information, see [“Vaulting Media Not Created by NetBackup”](#) on page 129.

Disk Only Source of Backups

If **Disk Only** is specified on the **Duplication** tab, an image that has no disk copy will not be duplicated even if a copy of that image exists on removable media and was selected during the Choose Backups step.

For more information about image selection, see [“The List of Images to be Vaulted”](#) on page 104.

The Scope of the Source Volume Group

The **Source Volume Group** on the **Choose Backups** tab spans all steps of a Vault profile (most notably Duplication and Eject). However, if you do not duplicate images, you do not have to specify a source volume group (the **Source Volume Group** field is ignored). Conversely, the **Source of Backups...** field on the **Duplication** tab applies only to the Duplication step.

Even if no images are selected for duplication, images may still be ejected if they are in the **Source Volume Group** and in an off-site volume pool specified on the profile **Eject** tab.



Using the Menu User Interfaces

10

Usually, you will use the NetBackup Administration Console to configure and run Vault. Vault also includes the following two menu user interfaces (MUIs) that you can use in a terminal window:

- ◆ The Vault Administration interface, which lets you configure Vault. It provides the same functionality as Vault Management in the NetBackup Administration Console.
- ◆ The Vault Operator Menu interface, which provides a way to eject media and print reports for one or more Vault sessions. You use the `vltopmenu` command to start the Vault Operator Menu interface.

This chapter covers the following topics:

- ◆ [Using the Vault Administration Interface](#)
- ◆ [Using the Vault Operator Menu Interface](#)
- ◆ [Changes in `vmadm` for Vault](#)
- ◆ [Changes in `bpdbjobs` for Vault](#)

Using the Vault Administration Interface

The Vault Administration interface (available on UNIX systems only) lets you configure and run Vault from a text-based menu. You can perform the same actions in the Vault Administration menu as in the NetBackup Administration Console.

You can use the Vault Administration interface from any character-based terminal (or terminal emulation window) that has a `termcap` or `terminfo` definition. Use the `vltaadm` command to start the Vault Administration interface, and run the `vltaadm` command only from the UNIX system on which the NetBackup master server resides. You must have root privileges to run the `vltaadm` command.

The `vltaadm` command and interface is available on UNIX systems only.



Note When you create or modify Vault configuration information, run only one instance of the NetBackup Administration Console or the Vault Administration interface. Using multiple instances at the same time may cause configuration information to be overwritten.

The `vltadm` command resides in the following directory:

`/usr/opensv/netbackup/bin`

When you invoke the `vltadm` command, the following menu appears in the terminal window:

```
Vault Administration
-----
      Robot Name:  <ALL>
      Vault Name:  <ALL>
      Profile Name: <ALL>

r)  Browse all configured robots
v)  Browse all configured vaults for selected robot
p)  Browse all configured profiles for selected vault

n)  Robot management...
t)  Vaults for selected robot...
f)  Profiles for selected vault...

c)  Copy selected profile...
s)  Start session for selected profile...

a)  Vault properties...
h)  Help
q)  Quit (without saving)
x)  Save and Exit          ENTER CHOICE:
```

The robot name, vault name, and profile name default to ALL. To browse through specific robots, vaults, or profiles already configured in Vault, press **r**, **v**, or **p**; the robot, vault, or profile names at the top of the menu will change. When the correct robot, vault, or profile is displayed, type the letter of the action you want to perform.

The criteria you can configure in the Vault Administration interface are described in [“Configuring Vault”](#) on page 49.

For help on the currently displayed menu, select the help option on that menu. Help includes a tutorial for learning and using the Vault Administration interface.

Using the Vault Operator Menu Interface

The Vault Operator Menu interface lets an authorized user eject and inject tapes and print reports for one or more Vault sessions (an authorized user is one who can invoke the `vltopmenu` command). The `vltopmenu` command, which starts the Vault Operator Menu, resides in the following directory:

UNIX: `/usr/opensv/netbackup/bin`

Windows: `install_path\NetBackup\bin`

When you invoke the `vltopmenu` command, the following menu appears in the terminal window:

```

                                NetBackup Vault Operator Menu

Current Profile: None
Current Session: 0
Current Report Destinations - Print command: /usr/ucb/lpr
                                Email:
                                Directory:

p) Select Profile                m) Modify the Report Destinations...
u) Profile Up                    r) Run Reports for This Session
d) Profile Down                  v) Run Individual Reports...
s) Select Session

                                cr) Consolidate All Reports
i) Inject Media into Robot        ce) Consolidate All Ejects
e) Eject Media for This Session   re) Consolidate All Reports and Ejects

                                c) Container Management...

q) Quit
Selection-->
```

Upon startup, the menu displays the current profile, session, and report destinations.

You can view the results of each operation in a log file. The name and location of the log file is located at the end of the output for each command. For example, if you choose **Eject Media for This Session** on a UNIX system, the output is similar to the following:

```

vlteject Started
vlteject Completed
The results of this operation have been logged in the following file:
/usr/opensv/netbackup/vault/sessions/vlteject_status/details.log.timestamp
```

Note Do not run another session for this vault while using this menu.



Changes in vmadm for Vault

The command line media management utility, `vmadm`, manages volumes, manages volume pools, manages barcode rules, and inventories robots controlled by the Media Manager volume daemon (`vmd`). The `vmadm` utility has a character-based user interface and can be used from any terminal. You must have root privileges to execute this utility.

When Vault is installed, several fields are added to the Volume Configuration display. You can modify these fields by using the Special Actions menu in `vmadm`.

Additions to Volume Configuration

When you display the full Volume Configuration, Vault specific information is displayed. The following is an example of the Vault specific information:

```
volume group: TL8-0
vault name: V1
vault sent date: Wed Dec 02 09:34:01 2000
vault return date: Tue Feb 17 09:34:01 2001
vault slot: 546
vault session ID: 37
created: Mon Nov 29 08:29:03 2000
```

Changes to the Special Actions Menu

You can modify several Vault options by using the Special Actions menu in `vmadm`, described in the following subsections.

Change Vault Name for Volumes

You can set, clear, or change the name of the vault in which the volume resides. This field is used by Vault to determine where the volume is located when it is at an off-site location.

▼ To modify the name of the vault

1. On the main menu, choose **s** for Special Actions.
2. Choose **a** for Change Vault Parameters for Volumes.
3. Choose **n** for Change Vault Name for Volumes.
4. Enter the new vault name (25 character maximum). Enter a hyphen (-) to clear the field.

5. You will be prompted for the media IDs for which you want this vault name applied. The prompt will repeat until you press the Enter key without entering a media ID. Click the ESC key to cancel the action.

Change Date Volumes are Sent to Vault

You can set, clear, or change the date a volume is sent to the off-site vault. This field is used by Vault to record when a volume was sent to the off-site vault location. You can modify this date for a single volume or for multiple volumes.

▼ To modify the Vault Sent Date of the vault

1. On the main menu, choose **s** for Special Actions.
2. Choose **a** for Change Vault Parameters for Volumes.
3. Choose **d** for Change Date Volumes are Sent to Vault.
4. Enter the new date the volume was sent off-site. Enter a zero (0) to clear the field.
5. You will be prompted for the date on which the volumes were sent to off-site storage. The prompt will repeat until you press the Enter key without entering a date. Press the ESC key to cancel the action.

Change Date Volumes Return from Vault

You can set, clear, or change the date a volume returns from the off-site vault. This field is used by Vault to record when a volume is requested to return from the off-site vault location. You can modify this date for a single volume or for multiple volumes.

▼ To modify the Vault Return Date of the vault

1. On the main menu, choose **s** for Special Actions.
2. Choose **a** for Change Vault Parameters for Volumes.
3. Choose **r** for Change Date Volumes Return from Vault.
4. Enter the new date the volume is requested to return from the off-site vault. Enter a zero (0) to clear the field.
5. You will be prompted for the return dates on which the volumes should be recalled from off-site storage. The prompt will repeat until you press the Enter key without entering a return date. Press the ESC key to cancel the action.



Change Vault Slot for Volumes

You can set, clear, or change the slot that the volume is contained in at the off-site location. This field is used by Vault to determine in what slot the volume is located in the off-site vault. You can modify the slot for a single volume or for multiple volumes.

▼ To modify the Slot ID for a volume

1. On the main menu, choose **s** for Special Actions.
2. Choose **a** for Change Vault Parameters for Volumes.
3. Choose **s** for Change Vault Slot for Volumes.
4. Enter the new vault slot ID for the volume. Enter a zero (0) to clear the field.
5. You will be prompted for the slot IDs to assign to the volumes. The prompt will repeat until you press the Enter key without entering a slot ID. Press the ESC key to cancel the action.

Change Vault Session ID for Volumes

You can set, clear, or change the session ID in which a volume was processed. This field is used by Vault to determine in what vault session a volume was processed. You can modify the session ID for a single volume or for multiple volumes.

▼ To modify the Session ID for a volume

1. On the main menu, choose **s** for Special Actions.
2. Choose **a** for Change Vault Parameters for Volumes.
3. Choose **i** for Change Vault Session ID for Volumes.
4. Enter the new session ID for the volume. Enter a zero (0) to clear the field.
5. You will be prompted for the session IDs to assign to the volumes. The prompt will repeat until you press the Enter key without entering a session ID. Press the ESC key to cancel the action.

Changes to Display Options

By default, information about all configured volumes is displayed in `vmadm`. You can set a filter option to limit the display to the volumes in a specific container.

▼ To display information about volumes in a container

1. On the main menu, choose **p** for Print Information about Volumes.
2. Choose **f** for Filter.
3. Choose **8** for VAULT CONTAINER ID.
4. Enter the container ID or a dash (-) to clear the container ID field and not filter on a container ID.

Changes in bpdjobs for Vault

The NetBackup activity monitor utility, `bpdjobs`, displays Vault jobs by default. If you invoke the `bpdjobs` command and use the `-vault` option, the output includes the following Vault specific fields.

Vault Fields in `bpdjobs` Output

| Field | Description |
|----------------|---|
| Robot | The name of the robot with which the vault is associated. |
| Vault | The name of the vault under which the session is running. |
| Profile | The name of the profile that holds the configuration information for the vault session. |
| Session ID | The session ID, a unique numeric value, for the vault job. Session ID assignment starts at 1 the first time a vault job is run after vault has been installed. The value is incremented by 1 every time a new vault job runs. |
| Tapes to Eject | The number of tapes to be ejected for a vault session. If the profile is configured for deferred eject, the tapes may not be ejected yet. |
| Operation | For Vault jobs, the field contains one of the following values. These values progress from the first value to the last as the Vault job progresses: <ul style="list-style-type: none"> ♦ Choosing Images ♦ Duplicating Images ♦ Choosing Media ♦ Catalog Backup ♦ Eject and Report ♦ Done |



If a Vault job completes successfully (with exit status = 0), the State field and the Operation field both contain the value Done. If a vault job fails, the Operation field contains the operation occurring at the time the job failed.

For information about specific problems or areas that can cause problems, see the following:

- ◆ [“Printing Problems”](#) on page 197
- ◆ [“Errors Returned by the Vault Session”](#) on page 198
- ◆ [“No Media Are Ejected”](#) on page 198
- ◆ [“Media is Missing in Robot”](#) on page 198
- ◆ [“Bad or Missing Duplicate Tape”](#) on page 199
- ◆ [“Tape Drive or Robot Offline”](#) on page 200
- ◆ [“No Duplicate Progress Message”](#) on page 200
- ◆ [“Ejecting Tapes While in Use”](#) on page 201
- ◆ [“Tapes Ejected to the MAP are Returned to Robot”](#) on page 201
- ◆ [“Unexpired Tapes Were Injected into the Robot”](#) on page 201
- ◆ [“Vault Session Locking”](#) on page 202
- ◆ [“Debug Logs”](#) on page 202

Printing Problems

Problems with printing reports that appear to be Vault problems often are problems with the print command configured on the profile **Reports** tab. Therefore, you should test print commands from a command line on the server on which Vault is installed to ensure that they work correctly.

In some rare cases with Microsoft Windows, the print command will work correctly when tested from a command prompt but will not work when configured on the profile **Reports** tab. The issue may be with how Windows calls the print command. If you experience such a problem, from a command prompt on the master server on which Vault is installed enter the following command (use the appropriate server and printer names):

```
NET USE lpt1 \\servername\printername PERSISTENT:YES
```



This problem can also occur in mixed environments of UNIX and Windows.

Errors Returned by the Vault Session

Every Vault session writes a detailed error status to `stderr`.

- ◆ If the error generated by the Vault session is less than or equal to 255, it will return the actual error code. Error codes less than or equal to 255 (except 252) map to standard NetBackup error codes and are documented in the *NetBackup Troubleshooting Guide*.
- ◆ If the Vault session fails with an error code greater than 255, it will return error code 252 and the actual error code will be written to `stderr`. This is because codes greater than 255 are called NetBackup Extended Error Codes and are not supported by all operating systems.

The format of the error text written to `stderr` is:

EXIT status = *error code*

For detailed information on status codes, see the *NetBackup Troubleshooting Guide for UNIX and Windows*.

No Media Are Ejected

If no media are ejected, it may be because of the following:

- ◆ All images have already been vaulted, so no images were selected. Vault determines that a backup image has already been vaulted if a copy of the image already is on a volume in an off-site volume group.
- ◆ The media to be vaulted are in a volume group other than the robotic volume group specified for the vault to which the profile belongs.

Media is Missing in Robot

Duplication may fail if NetBackup does not know that a requested piece of media is in the robot. For example, a tape may have been moved to the off-site volume group inadvertently even though it remains in the robot. To compare the tapes actually stored in the robot with the Media Manager database, use the NetBackup Administration Console **Inventory Robot** option.

If the tape is in the robot, use the NetBackup Administration Console to move the tape to the robotic volume group.

If the tape is not found, you should delete it from the NetBackup system. If the tape is missing yet is assigned and has valid duplicate images, you will need to use the command `bpexpdate`, which is documented in the *NetBackup System Administrator's Guide*, to expire the images before you delete the tape from Media Manager.

Bad or Missing Duplicate Tape

If a duplicate tape is lost or damaged, you can reduplicate the images that were on the tape if the primary backup images still reside in the robot.

▼ To reduplicate images

1. Determine which images were on the tape by running the `bpimmedia` command.
The `bpimmedia` command scans the entire NetBackup image catalog, so it may take a few minutes depending on the size of that catalog. Save the output because you will need to verify that the correct images were reduplicated.
2. Expire the lost or damaged duplicate tapes by using the `bpexpdate` command.
3. Determine when the images were created by using the `bpimagelist` command.
4. Create a profile that has the same criteria as the profile that created the missing duplicate tape except for the following:
 - ◆ Specify policy names only for the policy names used to create the images on the missing tape.
 - ◆ Set the time window so the profile selects the images on the missing tape. For example, if the original backups were made 30 days ago, set the time window between 32 and 28 days ago.
5. Run the profile by selecting it in the Administration Console and then selecting **Actions > Start Session**.

Ensure that no other Vault sessions are running before running this new profile.

Before duplicating images, you can verify that the correct images are selected by previewing the session. For more information, see [“Previewing a Vault Session”](#) on page 100.



Tape Drive or Robot Offline

If you have a problem with ACSLS drives going offline, you should try to configure the drives in an UP state or reset the drive. If drives persistently go offline, duplication may hang.

Also, if the tape drives are listed as AVR control in the Administration Console Media and Device Management node, there may be a problem with the robotics control. All drives should be listed as robotically controlled (that is, TLD, ACS, and so on), but they will be converted to AVR control if a problem occurs with the robot. To diagnose the problem, examine the system logs (for example, `/var/adm/messages` on UNIX systems) for error messages. You can also use the robotic test utilities (such as `robtest`) to further debug the problem.

No Duplicate Progress Message

If you see a message similar to the following in the `Vault detail.log`, the Vault process has not received any new information from the `bpduplicate` process within the time frame specified (in this example, 30 minutes):

```
bpduplicate_progress_logname: no activity in 30 minutes
```

bpduplicate_progress_logname is the progress log that `bpduplicate` creates as it runs the duplication for Vault. This file resides in the following directory:

UNIX:

```
/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx/duplicate.log.n
```

Windows:

```
install_path\NetBackup\vault\sessions\vault_name\sidxxx\duplicate.log.n
```

vault_name is the name of the vault used for the session, *xxx* is the unique session ID, and *n* is the number of the instance of the `bpduplicate` command (1 for the first instance, 2 for the second, and so on).

This message does not necessarily indicate an error has occurred. If the image that is currently being duplicated is very large (for example, several gigabytes), this message is displayed only for informational purposes. To determine if a problem exists, you can determine the size of the current image. First examine the last few lines of the `details.log` file to determine backup image ID. Then execute the `bpimagelist` command and specify the image ID, as in the following example:

UNIX: `bpimagelist -L -backupid server2_0897273363`

Windows: `bpimagelist.exe -L -backupid server2_0897273363`



The output of this command will show you various statistics about this backup image, including the number of kilobytes written during this backup. If the number is relatively small, there may be a problem with the duplication process. Sometimes this delay is caused by a media mount (which normally does not occur in robotic devices during duplication), by hardware problems, or by the media being in use. Examine the Activity Monitor to determine if there are any hardware problems and also check the system logs. If the backup image is very large, then this message should be regarded as informational.

Ejecting Tapes While in Use

If Vault is configured to eject original media, it is possible that a piece of media could be in use during the eject process (for example, for a restore or a media verify procedure.) In this case, an error message will be generated by Media Manager. A similar error may be generated by non-Media Manager controlled robots if a piece of media is currently in use.

If you receive one of these errors, we recommend that you use the Vault Operator Menu (`vltopmenu`) to re-eject the media after the media is no longer in use. You may receive additional error messages because the rest of the media for the scheduled job has already been ejected.

Tapes Ejected to the MAP are Returned to Robot

If media is not removed from the robot's media access port (MAP) and a timeout condition occurs, the media is returned to (injected into) the library slots in the robot. If this occurs, the Vault reports will not accurately reflect the status of the media.

To recover, you should use the Vault Operator Menu (`vltopmenu`) or `vlteject` to eject the media that was not removed from the library and generate the reports.

Unexpired Tapes Were Injected into the Robot

If Vault tapes that have not expired have been injected back into a robot, you can revault them manually by doing the following:

1. Eject the media, as follows:
 - a. In the NetBackup Administration Console, select the robot into which the media was injected (**Media and Device Management > Robots**).
 - b. Select the media ID(s) you want to eject.
 - c. Select the **Eject Volumes from Robot....** operation on the **Actions** menu.



2. Transfer the media to the off-site volume group by doing the following:
 - a. Select the media ID(s).
 - b. Select **Actions > Change Volume Group**.
 - c. Choose the appropriate off-site volume group from the **New Volume Group Name** drop down menu.
3. Return the media to your vault vendor so that all backups on that media will be available for future disaster recovery.
4. Run the *Recovery* report to ensure that the media is available for future disaster recovery operations.

Alternatively, you can use the `vmchange` command to eject the media and transfer it to the off-site volume group.

Vault Session Locking

For any vault, only one profile can run at a time. If a second vault session is started while a first is already active for the vault name, the second vault session will terminate immediately, and the following message will be displayed:

```
A session is already running for this vault
```

The vault lock file is located in the directory for the specified vault:

UNIX:

```
/usr/opensv/netbackup/vault/sessions/vault_name/vault.lock
```

Windows:

```
install_path\NetBackup\vault\sessions\vault_name\vault.lock
```

Debug Logs

Vault debug logs are created in a `vault` directory in the the standard NetBackup debug logging path. You must create the `vault` directory so that daily log files are generated; if the directory does not exist, the log files will not be created. The following are the daily Vault debug logs:

UNIX: `/usr/opensv/netbackup/logs/vault/log.mmdyy`

Windows: `install_path\NetBackup\logs\vault\mmdyy.log`



The amount of information logged and how long it is retained is controlled by a NetBackup configuration parameter.

VERITAS recommends that you use a debug level of 5 when you generate logs that you send to VERITAS for troubleshooting purposes. You can set the debug level to 5 globally or you can use the *-verbose* option on the `vltrun` command in the Vault policy that invokes the Vault job.

NetBackup Debug Logs

If you want to log activity for NetBackup processes, you must create a subdirectory for the specific NetBackup software you are interested in monitoring. For example, if you want to monitor the `bptm` process to see the tape manager output, you must create a `bptm` directory on the server where the process runs.

For more information about the NetBackup debug logs, see the *NetBackup Troubleshooting Guide for UNIX and Windows*.

Setting the Duration and Level of Debug Logs

You can use the NetBackup Administration console to set the length of time NetBackup retains debug logs and the level of information contained in them. The setting affects all log files generated by NetBackup.

▼ To set the duration and level for log files:

1. In the NetBackup Administration Console, expand **NetBackup Management**.
2. Expand **Host Properties**.
3. Select **Master Server**.
4. In the right pane, select the master server and **Actions > Properties**.
5. Select **Global NetBackup Attributes**.
6. Enter the length of time to retain the NetBackup log files. This setting applies to all NetBackup logs, including but not limited to, the Vault logs.

In the NetBackup Administration Console for UNIX, the field name is **Keep Logs For**.

In the NetBackup Administration Console for Windows, the field name is **Duration to Retain Logs**.

7. Select the **Logging** tab.



8. Enter the logging level.

The logging level determines how much information is displayed in the log. A level of 0 will display the minimum amount of information; a level of 5 will display the maximum.

In the NetBackup Administration Console for UNIX, select a **Vault Logging Level**. The logging level corresponds to the `bp.conf` entry `VAULT_VERBOSE = level`.

In the NetBackup Administration Console for Windows, select a **Vault Logging level**.

Logs To Accompany Problem Reports

To troubleshoot problems, VERITAS Customer Service requires a set of log files produced by NetBackup and Vault. In most circumstances, you will have to provide log files from the following NetBackup processes:

- ◆ `admin` (on the Master Server); administrative commands process
- ◆ `bpbrmvlt` (on the Master Server); Vault job scheduler (similar to `bpbrm`)
- ◆ `bpcd` (on the Master Server); NetBackup client daemon manager
- ◆ `bpsched` (on the Master Server); NetBackup backup scheduler
- ◆ `bptm` (on the Media Server); NetBackup tape manager
- ◆ `vault` (on the Master Server); Vault process

Session log files also are useful for troubleshooting problems, and you should include the appropriate session log files with problem reports you send to VERITAS.

If you use the `vlteject` command or the Vault Operator Menu (`vltopmenu`) to perform consolidated ejects and reports, the following log file may also be useful:

UNIX: `/usr/opensv/netbackup/vault/sessions/vlteject.mstr`

Windows: `install_path\NetBackup\vault\sessions\vlteject.mstr`

Recovering from Disasters



This section provides information about recovering data, using NetBackup and Vault, when you have to recall your media from your off-site storage location. It also provides general information about preparing for a disaster recovery situation. For more information, see the following:

- ◆ [“Introduction”](#) on page 205
- ◆ [“Disaster Recovery in the NetBackup Vault Context”](#) on page 209
- ◆ [“Preparing for Recovery”](#) on page 209
- ◆ [“Recovering NetBackup”](#) on page 211
- ◆ [“Recovering Data”](#) on page 212
- ◆ [“Recovering to a Specific Point in Time”](#) on page 215

For information about recovering the NetBackup application, see the “Disaster Recovery” section in the *NetBackup Troubleshooting Guide for UNIX and Windows*.

Introduction

Data backup is essential to any data protection strategy, especially a strategy that is expected to assist in disaster recovery. Regularly backing up data and then being able to restore that data within a specified time frame are important components of recovery. Regardless of any other recovery provisions, backup protects against data loss from complete system failure. And off-site storage of backup images protects against damage to your on-site media or against a disaster that damages or destroys your facility or site.

To perform recovery successfully, the data must be tracked to know at what point in time it was backed up, which allows your organization to assess the information that cannot be recovered. Your data backup schedules should be configured to allow your organization to achieve its recovery point objective (RPO), which is the point in time before which you cannot accept lost data. If your organization can accept one day’s data loss, your backup schedule should be at least daily so you can achieve an RPO of one day before any disaster.



Your organization also may have a recovery time objective (RTO), which is the expected recovery time or how long it will take to recover. Recovery time is a function of the type of disaster and of the methods used for recovery. You may have multiple RTOs, depending on which services your organization must recover and when.

High availability technologies can make the recovery point very close or even identical to the point of failure or disaster, and they also can provide very short recovery times. However, the closer to the failure that you place your RTO and RPO, the more expensive it becomes to build and maintain the systems required to achieve recovery. Your analysis of the costs and benefits of various recovery strategies should be part of your organization's recovery planning. Understanding disaster recovery planning allows you to place Vault and tape-based backups stored off-site in the proper context within your disaster recovery objectives.

Definition of Disaster

For an organization, a disaster is an unplanned event that interrupts its ability to function. Usually, the event affects the delivery of critical business functions and results in a loss of data. The following are generally recognized as the types of disasters possible:

- ◆ Technological disasters result in shortcomings in performance, availability, capacity, and accessibility of your IT infrastructures. Technological disasters include computer or Internet crimes, computer viruses, power failures, network or telecommunication failures, hardware or software failures, and so on.
- ◆ Human disasters are caused by people, including accidents, explosions, fires, riots, terrorist activities, and so on.
- ◆ Natural disasters are caused by nature, including hurricanes, tornadoes, earthquakes, floods, and so on.

The impact of a disaster often depends on the scale and timing of the event. Although a disaster is an event that is beyond your control, you can control the way in which your organization reacts to a disaster. By planning and preparing for a disastrous event, you can minimize the effect of the disaster.

Definition of Disaster Recovery

Disaster recovery is the process of responding to an interruption in the services your organization uses to operate. Disaster recovery usually is focused on information, network, and telecommunication services, often at an alternative site and using one or more data-recovery methods.

Disaster recovery is part of a larger topic called business recovery, which is restoring the actual capability for employees to perform their jobs. Business recovery includes logistic related items, such as telephones, office space, living arrangements for employees, and so

on. Business recovery itself is part of a larger topic called business continuity planning, which includes plans to manage the crisis to your organization, help resume normal business operations, and so on.

A resilient organization will use business continuity planning to help ensure that it can survive a disaster and resume operations at an acceptable level.

Definition of Disaster Recovery Plan

A disaster recovery plan is a plan to resume or recover a specific essential operation, function, or process of an organization. Although disaster recovery usually has been used to describe information technology and telecommunication services recovery, other services an organization uses to conduct operations can and should be considered part of a plan. For example, an organization's people also are subject to the effects of a disaster and planning should include the impact on them and the resources necessary to help them recover so they can perform their duties.

By planning how your company will respond in the event of a disaster, you ensure that your company can:

- ◆ Protect critical data
- ◆ Minimize the impact of a disaster
- ◆ Use resources most effectively
- ◆ Maintain business continuity

Recovery Priorities

Your organization must decide between recovery cost (the infrastructure and testing) and the level of functionality that must be recovered. You may choose to recover only the most critical business functions immediately and then recover other functions later. Although all functions of an organization should be valuable and necessary for the organization to operate, it may be acceptable to operate at a reduced level for a specific period of time. The longer your organization can operate without a function, the easier and less expensive it becomes to recover that function. Therefore, given the higher cost of rapid recovery, only those functions required for immediate operation need to be recovered quickly. Delaying recovery of some functions can be a good business decision.

Developing a Disaster Recovery Plan

Developing a disaster recovery plan usually begins with an impact analysis that identifies the functions an organization requires to operate and determines how long each function can be unavailable until it affects the organization to an unacceptable extent.



Understanding the impact of disaster will help you identify the objectives for the recovery plan. The following are examples of objectives that may be in a disaster recovery plan:

- ◆ Ensure service to customers by making critical resources available
- ◆ Minimize economic loss
- ◆ Secure company assets
- ◆ Minimize decision making during the recovery process
- ◆ Reduce reliance on key individuals
- ◆ Ensure a safe and orderly recovery within predetermined time period
- ◆ Maintain a sense of security and organizational stability

The priority you assign your objectives depends on the needs of your organization. By setting clear, prioritized objectives for your disaster recovery plan, you can reduce your organization's exposure to risks and ensure that your critical systems and networks are available during disruptions.

Two approaches can be used to create disaster recovery plans:

- ◆ A general plan that is used any time a disaster occurs. A general plan should be flexible and is often impact driven rather than disaster driven (that is, based on the impact to your organization rather than the type of disaster). A general plan usually is based on assumptions that define the scope of each impact in the plan. A general plan is easy to maintain and convenient; however, because it may require that some decisions be made at the time of disaster (such as assessing the type of impact and determining the response), the beginning of recovery may be delayed.
- ◆ Multiple smaller plans, each used for a specific disaster that your organization has determined is most likely to occur. For example, individual plans often are created for power outages, network outages, fires, floods, and so on. Individual disaster-specific plans are easier to create than a general plan. It is often clear which plan should be used, so fewer decisions are required at the beginning of recovery, which can result in quicker recovery. However, which plan to use may not always be clear (for example, if a fire causes a power outage), and if a disaster occurs for which a plan does not exist, recovery may be slow to begin and difficult to achieve.

A disaster recovery plan should be easy to follow and not require interpretation. Do not include unnecessary detail. If the plan is implemented, it will be in a time of high stress and pressure to perform; therefore, it should be simple, specific, and well tested.

You should publicize your disaster recovery plan within your organization so that everyone knows about it, understands how it works, and understands the reasoning behind the decisions in the plan.

Testing a Disaster Recovery Plan

Developing a disaster recovery plan is a waste of time and resources if it is not tested regularly, thoroughly, and frequently (*frequently* depends on any changes in your organization's functions and environment). The goal of disaster recovery testing is not to pass but to find out what does not work. Your tests should be designed to find problems because it is better to find them during a test than during an actual recovery situation.

Testing can be as simple as calling all of the phone numbers in an emergency notification list, which helps verify that everyone can be reached when needed. Or testing can be as complex as actually conducting operations at a recovery site to ensure that everything works correctly. Between those extremes, variations include walkthroughs, during which everyone involved in the recovery process discusses their roles in a moderated recovery scenario, and simulations that invoke the recovery plan but use simulated data. Perhaps using a combination of testing scenarios to test specific parts of the plan also can be effective.

Disaster Recovery in the NetBackup Vault Context

In the NetBackup Vault context, disaster recovery means restoring the NetBackup application (master server, media servers, and clients) and then restoring the data that is stored at the off-site storage facility.

If you use NetBackup and Vault to backup your applications and store removable media at a secure off-site location, you also can perform application recovery so you do not have to reinstall your applications.

Preparing for Recovery

Note Effective disaster recovery procedures are specific to an environment and provide detailed information about everything that should be accomplished to prepare for disaster and to recover after disaster occurs. VERITAS provides general disaster recovery information that is intended as a model only; you must evaluate the information and then develop your own disaster recovery plans and procedures.

Recovering data can be a difficult and time consuming process. The success of recovery often depends on how well you prepare for disaster. Your preparations for disaster and what you have to accomplish during a recovery depends on your recovery systems. For example, if your recovery site and systems are already operational and have NetBackup and Vault installed, you do not have to protect the NetBackup installation media and the license keys and install NetBackup during the recovery process — you only have to



recover the NetBackup catalogs and data. Conversely, if your recovery systems do not have NetBackup and Vault installed and configured, you have to prepare for that and accomplish it during recovery.

You should do the following to prepare for recovery using NetBackup and Vault; you may not have to do some of the items listed, and you may have to do more than what is listed:

- ◆ Develop a disaster recovery plan.
- ◆ Test the disaster recovery plan.
- ◆ Back up and vault data regularly. In addition to backing up files on a regular basis, it is important to select the correct files to back up. You should back up all data that your organization's impact analysis determines is critical and store copies at a secure, off-site storage location.
- ◆ If you can recover to the same or identical hardware, back up and vault the applications that your organization's impact analysis determines are critical. You also should back up system files so you can quickly restore a system to normal operation:

- ◆ Include all operating system files in your backups. For Microsoft Windows systems, the Windows system directories include the registry, without which it is impossible to restore a system to its original configuration. If you are using a NetBackup exclude list for a client, do not specify any Windows system files in that list.

Restoring operating system files is not helpful if you are recovering data to a different system at your original or disaster recovery site. You can back up those files, but then not restore them if you are recovering to a different system or site.

- ◆ Back up executable and other files for applications you need to conduct critical operations. You may want to save money by using fewer tape volumes, but backing up critical applications ensures that you can restore them to their exact configurations. For example, if you have applied software updates or patches, restoring from a backup eliminates the need to reapply them, reducing recovery time.
- ◆ Every time you vault media, store the Recovery Report securely. The same disaster that destroys your site can destroy your Recovery Report. You will need the Recovery Report to identify the media you need to recall from off-site storage. Your vault vendor may allow you to vault your Recovery Report. If you have a recovery site equipped with computers, e-mail the Recovery Report to that site.
- ◆ Record and protect the names of the policies that are used to backup the data you want to recover. The Recovery Report is organized by policy, so you need to know which policies are used so you can identify the media you need to recover.

- ◆ Record and protect the names of the off-site volume groups for the data you want to recover. Those names are used during the recovery process. Alternatively, you can obtain the off-site volume group names after you restore the NetBackup catalog (because the catalog includes the Vault configuration).
- ◆ Document the commands and options that you need to recover data. For example, the `bpchangeprimary` command is used to promote the vaulted images to primary images so that you can restore from them. So you should have a record of the commands and options that you need during the recovery process.
- ◆ Protect the NetBackup and Vault installation media. You need the media so you can install NetBackup and Vault on the recovery system if it is not already installed.
- ◆ Record and protect the license keys for NetBackup and Vault. You need them for NetBackup and Vault on the recovery system if you have to install NetBackup. You can use temporary license keys if necessary.
- ◆ Protect the installation media and record the license keys for any other VERITAS products that must be installed on the recovery systems. For example, if you use the VERITAS File System™ and VERITAS Volume Manager™ on the recovery systems, you will need their license keys when you install those products.
- ◆ Protect the installation media for the operating system and other applications required to run the systems you are using for recovery.
- ◆ Protect your DR plan. The same disaster that destroys your site can destroy your DR plan and recovery report. You should have copies stored so that they will be available when needed. Your vault vendor may allow you to vault a copy of the DR plan.

Recovering NetBackup

Information about recovering NetBackup master servers, media servers, and client systems after a disk failure is provided in the “Disaster Recovery” section of the *NetBackup Troubleshooting Guide for UNIX and Windows*. The procedures include re-installing NetBackup and may include re-installing the operating system if it is required.

You must first ensure that your computer and robotic hardware and any necessary network equipment is operational before you re-install NetBackup. You can then use an appropriate procedure in the *NetBackup Troubleshooting Guide for UNIX and Windows*.

After you re-install NetBackup, you may have to configure robots and drives (unless you re-install NetBackup on the original system, in which case the device configuration will be restored if you recover the NetBackup catalogs).



Recovering Data

Note Effective disaster recovery procedures are specific to an environment and provide detailed information about everything that should be accomplished to prepare for disaster and to recover after disaster occurs. VERITAS provides general disaster recovery information that is intended as a model only; you must evaluate the information and then develop your own disaster recovery plans and procedures.

Recovering data can be a difficult and time consuming process. The success of recovery often depends on how well you prepare for disaster and subsequent recovery.

The steps you have to perform to recover can depend on the configuration of your original system and robotic devices, your original NetBackup configuration, the configuration of your recovery system and robots, and the configuration of NetBackup on the recovery systems. Therefore, providing specific disaster recovery procedures for all situations is not possible; rather, these procedures are intended as general guidelines from which you can create your own procedures for recovering NetBackup and the data that was protected by transferring it off-site.

Although some detail is included about restoring the NetBackup catalogs to a recovery system, these procedures do not provide detail about every step of the recovery procedure.

Information in this section assumes the following:

- ◆ Primary backup images unavailable.
- ◆ NetBackup master and media server software, Vault software, client software, and devices installed and robots and drives configured on systems to which you are recovering data.
- ◆ NetBackup catalogs and data media not recalled from off-site storage.
- ◆ Recovery Report available.
- ◆ Off-site volume group to which the recovered images belong is known.
- ◆ Master and media server names of the original systems are known.

▼ To recall media and restore backup images

1. Using the Recovery Report, identify the current catalog backup media and the media required to restore the data.

The Recovery Report is organized by policy, so you should determine which policies were used to backup the data you want to recover.

2. Recall the appropriate catalog backup and data media from off site storage.

3. On the recovery system, rename the NetBackup device configuration files so you can restore the recovery system device configuration after you recover the catalogs.

By default, NetBackup catalog backups include the NetBackup device configuration files. If you recover the catalogs to a new system that has a device configuration different from the original system, the device configuration from the original system will replace that of the recovery system. Therefore, before recovering the catalog, rename the following files on the recovery system:

UNIX systems:

```
/usr/opensv/volmgr/database/ltidevs
/usr/opensv/volmgr/database/robotic_def
```

Windows systems:

```
install_path\Volmgr\database\ltidevs
install_path\Volmgr\database\robotic_def
```

Rename them so it is obvious what the original file names were; for example, rename them by giving them an extension such as .bak.

You will be directed to restore the recovery server's device configuration files in [step 5](#).

4. Recover the NetBackup catalogs using one of the procedures in the "Disaster Recovery" section of the *NetBackup Troubleshooting Guide for UNIX and Windows*.

During catalog recovery, do *not* restart the NetBackup daemons (UNIX) or services (Windows) as directed in the "Disaster Recovery" section procedures.

If you started the NetBackup Administration Console during catalog recovery, it can remain running.

5. On the recovery server, copy the files you created in [step 3](#) and give them the names of the original files:

UNIX systems:

```
/usr/opensv/volmgr/database/ltidevs
/usr/opensv/volmgr/database/robotic_def
```

Windows systems:

```
install_path\Volmgr\database\ltidevs
install_path\Volmgr\database\robotic_def
```

Renaming these files restores the device configuration for NetBackup on the recovery system.



6. If the master and media server names on the recovery system are different than the original system, change the server names in the NetBackup catalogs by using the `bpimage` command. The `bpimage` command and options required are as follows:

```
bpimage -newserver recovery_system -oldserver original_system
```

You can also use the `bpimage` command if the old system had separate media servers and the recovery system has a combined master and media server. Use the name of the combined master/media server for the argument to the `-newserver` option.

7. Start the NetBackup daemons (UNIX) or services (Windows) as directed in the “Disaster Recovery” section procedures.
8. Change the images to be recovered to primary (NetBackup restores from the primary image).

To change a large number of images to primary, you can use `bpchangeprimary -group` option to specify all images in a specific off-site volume group. For information about the `bpchangeprimary` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* guide.

9. If the media is not suspended or frozen, suspend the media.

Use the `bpmedia` command to suspend the media. Suspending the media prevents NetBackup from writing backup images to that media.

10. If the NetBackup Administration Console is not running, start it.

11. Inject the media into the robot. For procedures, see “[Injecting Media](#)” on page 112.

Injecting the media moves it into the robot and also changes the off-site volume group attribute of the media to robotic volume group so NetBackup knows that the volumes are in the robot and ready for restore operations.

12. Using the Backup, Archive, and Restore interface, restore the data. For procedures, see the *NetBackup User's Guide for UNIX* or *NetBackup User's Guide for Microsoft Windows*.
13. After restoring all the data, revault the media. For procedures, see “[Revaulting Unexpired Media](#)” on page 126.

Recovering to a Specific Point in Time

If your data center or computing environment requires recovery to a specific point in time (not just to the most recent valid backups), you can set up a process that will ensure that you can recover both the NetBackup catalog and the data for that specific time. You should retain the corresponding catalog backups for the same length of time as the corresponding data backups.

The following high-level information is intended as an overview of how to archive the catalog and data so you can recover to a specific point in time; detailed instructions for accomplishing all the tasks necessary are not included.

▼ To archive the catalog and data

1. Use your normal procedures to vault the data and NetBackup catalogs for that data.
2. Use the `bpmedia` command to freeze the data and catalog volumes that you want to retain.

Freezing the volumes prevents them from becoming unassigned and from appearing on the Picking List for Vault report. Do not recall the volumes from off-site storage when they expire.

3. Vault a printed copy of the Recovery Report for that specific point in time.

You will need the Recovery Report from the specific point in time so you can recall and restore the appropriate catalog and data volumes.

4. Optionally, remove the media IDs from the volume database.

This reduces the size of the database and improves performance. Depending on the number of volumes, maintaining the media IDs in the volume database may not degrade performance much.

▼ To recover the catalog and data

1. Retrieve the appropriate printed Recovery Report from off-site storage.
2. Using the Recovery Report, recall the appropriate catalog backup and data volumes from off site storage.
3. Stop the NetBackup daemons (UNIX) or services (Windows).
4. Restore the catalog that you recalled from off-site storage.

That version of the catalog contains information about the archived volumes and the images on them.



5. Use the `bpexpdate` command to reset the expiration date on the recalled volumes so they will not be expired. You can use the `-policy` option to change the expiration date for all media associated with a specific policy.
6. Change the images to be recovered to primary (NetBackup restores from the primary image).

To change a large number of images to primary, you can use `bpchangeprimary -group` option to specify all images in a specific off-site volume group. For information about the `bpchangeprimary` command, see the *NetBackup Commands for UNIX* or *NetBackup Commands for Windows* guide.

7. Restart the NetBackup daemons (UNIX) or services (Windows).
8. Restore the data.

The *NetBackup System Administrator's Guide, Volume I* includes an alternative procedure for archiving the catalog. That alternative procedure uses the *catarc* catalog archive policy to archive old data in the NetBackup catalog. You can then vault the archived catalog data or a copy of the archived catalog data. For more information, see "Catalog Archiving" in the *NetBackup System Administrator's Guide, Volume I*.

Vault's File and Directory Structure

For information about the files and directories created by Vault, see the following:

- ◆ [“UNIX Files and Directories”](#) on page 217
- ◆ [“Windows Files and Directories”](#) on page 224

UNIX Files and Directories

Vault is installed in `/usr/opensv/netbackup` on UNIX systems. The following table describes the files in each of the directories Vault creates and uses on UNIX systems. Also included are Vault files that reside in NetBackup directories. The files are either copied into the directories during the installation process or created as Vault sessions run. *The paths are specified in relation to the netbackup directory.*

Files and Directories for Vault in UNIX

| Directory, Program, or File | Purpose |
|--|---|
| <code>bin/bpbrmvl</code> | Process that kicks off vltrun from a scheduled or manual vault policy. |
| <code>bin/goodies/vlt_ejectlist_notify</code> | Script called by the vault session just before vault tapes are ejected. |
| <code>bin/goodies/vlt_end_notify</code> | Script called by the vault session just before it exits. |
| <code>bin/goodies/vlt_endeject_notify</code> | Script called by the vault session at the end of eject processing. |
| <code>bin/goodies/vlt_start_notify</code> | Script called by the vault session after it starts. |
| <code>bin/goodies/vlt_starteject_notify</code> | Script called by the vault session when the eject process starts. |



Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|-----------------------------|--|
| bin/vltadm | Utility which has a menu interface that an administrator can use to configure NetBackup Vault and monitor its operations. vltadm requires root (administrator) privileges. UNIX only. |
| bin/vltcontainers | Command used to add media logically to containers. |
| bin/vlteject | Command used to eject media from Vault sessions and run the reports selected in the profile. |
| bin/vltinject | Command used to inject media into a robot and update the Media Manager database. |
| bin/vltoffsitemedia | Command which allows the user to change the off-site parameters of a given piece of media. |
| bin/vltopmenu | Utility which allows the user to invoke a menu screen containing the various options that an operator of the NetBackup Vault feature can use. It allows the user to eject or inject media, print various reports individually or collectively, and consolidate reports and ejects across sessions. |
| bin/vltrun | Process that executes all the NetBackup commands used during a Vault session. |
| db/vault/retention_mappings | File that maps retention levels. Used to assign retention levels to duplicate images based on the retention level of the backup image. |
| db/vault/vault.xml | The vault configuration file. |
| help/vltadm | Contains help files for the Vault Administration (vltadm) interface. |
| vault | The main directory for Vault. It contains programs, session directories, and so on. |
| vault/sessions | A subdirectory that contains working session directories and log files. In a cluster environment, the sessions directory must reside on the shared disk. |

Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|--|--|
| <code>vault/sessions/cntrDB</code> | The database of information for media vaulted in containers. |
| <code>vault/sessions/sidxxx</code> | Subdirectory that contains working session subdirectories. Can be manually removed to reduce disk usage if necessary. |
| <code>vault/sessions/vault_name/session.last</code> | Counter to show current duplication session |
| <code>vault/sessions/vault_name/sidxxx/ allmedia_inventory.rpt</code> | The All Media Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ class.inventory</code> | Listing of all configured NetBackup policies; for use with recovery report. |
| <code>vault/sessions/vault_name/sidxxx/ container_inv.rpt</code> | The Container Inventory report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ detail.log</code> | Shows the output of every command that was executed during the session. |
| <code>vault/sessions/vault_name/sidxxx/ detailed_distlist_vault.rpt</code> | The Detailed Distribution List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |



Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|---|---|
| <code>vault/sessions/vault_name/sidxxx/ distlist_robot.rpt</code> | The Distribution List for Robot Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ distlist_vault.rpt</code> | The Distribution List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ duped.images</code> | List of images successfully duplicated during the session. |
| <code>vault/sessions/vault_name/sidxxx/ duplicate.log.nn</code> | Contains output from <code>bpduplicate</code> . |
| <code>vault/sessions/vault_name/sidxxx/ eject.list</code> | List of media to be ejected for the session. |
| <code>vault/sessions/vault_name/sidxxx/ image.list</code> | NetBackup image catalog information for each image duplicated. |
| <code>vault/sessions/vault_name/sidxxx/ image.list_suspend</code> | NetBackup image catalog information for each image whose media will be suspended. |
| <code>vault/sessions/vault_name/sidxxx/ imagefull.list</code> | A list of all images in the NetBackup image catalog. |
| <code>vault/sessions/vault_name/sidxxx/ lostmedia.rpt</code> | The Lost Media Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |

Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|--|--|
| <code>vault/sessions/vault_name/sidxxx/ media.list</code> | NetBackup media used for originals and duplicates. |
| <code>vault/sessions/vault_name/sidxxx/ media.list_suspend</code> | NetBackup media used for originals to be suspended. |
| <code>vault/sessions/vault_name/sidxxx/ nb_media.list</code> | Contains the number of images, size of images, and expiration dates for original and duplicated media. |
| <code>vault/sessions/vault_name/sidxxx/ nb_media.list_suspend</code> | Contains the number of images, size of images, and expiration dates on original media to be suspended. Used for reports. |
| <code>vault/sessions/vault_name/sidxxx/ nbudb.media</code> | Contains media IDs that were allocated for NetBackup database backups. |
| <code>vault/sessions/vault_name/sidxxx/ non_vaulted.rpt</code> | The Non-vaulted Images Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ notvaulted.list</code> | A list of images that were not vaulted. |
| <code>vault/sessions/vault_name/sidxxx/ offsite_inventory.rpt</code> | The Off-site Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ picklist_robot.rpt</code> | The Picking List for Robot report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |



Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|---|---|
| <code>vault/sessions/vault_name/sidxxx/ picklist_vault.rpt</code> | The Picking List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ preview.list</code> | Duplication preview output file. Shows images that are to be duplicated. |
| <code>vault/sessions/vault_name/sidxxx/ preview.list_suspend</code> | List of images for which media will be suspended. |
| <code>vault/sessions/vault_name/sidxxx/ rcvrimage.inventory</code> | NetBackup image catalog information for all policies between dates specified in the profile. For use with the recovery report. |
| <code>vault/sessions/vault_name/sidxxx/ recovery.rpt</code> | The Recovery Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/ returned_media.list</code> | Temporary file used for expiring recalled tapes from offsite vault. |
| <code>vault/sessions/vault_name/sidxxx/ robot.inventory</code> | Lists all the media currently residing in the robot (one media ID per line). |
| <code>vault/sessions/vault_name/sidxxx/ summary.log</code> | Concise view of detail.log listing major events during the session, such as how many images were copied. This log is appended for e-mail notification. |
| <code>vault/sessions/vault_name/sidxxx/summa ry_distlist_vault.rpt</code> | The Summary Distribution List report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |



Files and Directories for Vault in UNIX (continued)

| Directory, Program, or File | Purpose |
|--|---|
| <code>vault/sessions/vault_name/sidxxx/vault.err</code> | Error log for duplication of specific image. |
| <code>vault/sessions/vault_name/sidxxx/vault.err_suspend</code> | Error log for other administrative commands performed during suspend mode; this file should be checked in case of problems. |
| <code>vault/sessions/vault_name/sidxxx/vault.error.file</code> | Error log for other administrative commands performed by Vault; this file should be checked in case of problems. |
| <code>vault/sessions/vault_name/sidxxx/vault_inventory.rpt</code> | The Vault Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault/sessions/vault_name/sidxxx/volume.db.list</code> | Media Manager inventory for NetBackup database duplicate pool. |
| <code>vault/sessions/vault_name/sidxxx/volume.inventory_suspend</code> | Media Manager inventory for original media pools. |
| <code>vault/sessions/vault_name/sidxxx/volume.list</code> | Media Manager inventory for duplicate pool and NetBackup database duplicate pool. |
| <code>vault/sessions/vault_name/sidxxx/volume_full.list</code> | Media Manager inventory of all media. |



Windows Files and Directories

Vault is installed in the directory specified by *install_path\NetBackup* on Windows systems. The following table describes the files in each of the directories Vault creates and uses on Windows. Also included are Vault files that reside in NetBackup directories. The files are either copied into the directories during the installation process or created as Vault sessions run. *The paths are specified in relation to the NetBackup directory.*

Files and Directories for Vault in Windows

| Directory, Program, or File | Purpose |
|-----------------------------------|--|
| bin\bpbbrmvl | Process that kicks off vlrun from a scheduled or manual vault policy. |
| bin\goodies\vlt_ejectlist_notify | Script called by the vault session just before vault tapes are ejected. |
| bin\goodies\vlt_end_notify | Script called by the vault session just before it exits. |
| bin\goodies\vlt_endeject_notify | Script called by the vault session at the end of eject processing. |
| bin\goodies\vlt_start_notify | Script called by the vault session after it starts. |
| bin\goodies\vlt_starteject_notify | Script called by the vault session when the eject process starts. |
| bin\vltcontainers | Command used to add media logically to containers. |
| bin\vlteject | Command used to eject media from Vault sessions and run the reports selected in the profile. |
| bin\vlthinject | Command used to inject media into a robot and update the Media Manager database. |
| bin\vltoffsitemedia | Command which allows the user to change the off-site parameters of a given piece of media. |
| bin\vltopmenu | Utility which allows the user to invoke a menu screen containing the various options that an operator of the NetBackup Vault feature can use. It allows the user to eject or inject media, print various reports individually or collectively, and consolidate reports and ejects across sessions. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|--|--|
| <code>bin\vltrun</code> | Process that executes all the NetBackup commands used during a Vault session. |
| <code>db\vault\retention_mappings</code> | File that maps retention levels. Used to assign retention levels to duplicate images based on the retention level of the backup image. |
| <code>db\vault\vault.xml</code> | The vault configuration file. |
| <code>vault</code> | The main Vault directory. Contains programs, working directories, etc. |
| <code>vault\sessions</code> | A subdirectory containing working session directories and log files. In a cluster environment, the sessions directory must reside on the shared disk. |
| <code>vault\sessions\cntrDB</code> | The database of information for media vaulted in containers. |
| <code>vault\sessions\sidxxx</code> | Subdirectory containing working session subdirectories. Can be manually removed to reduce disk usage if necessary. |
| <code>vault\sessions\vault_name\session.last</code> | Counter to show current duplication session |
| <code>vault\sessions\vault_name\sidxxx\allmedia_inventory.rpt</code> | The All Media Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\class.inventory</code> | Listing of all configured NetBackup policies; for use with recovery report. |
| <code>vault\sessions\vault_name\sidxxx\container_inv.rpt</code> | The Container Inventory report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|---|---|
| <code>vault\sessions\vault_name\sidxxx\detail.log</code> | Shows the output of every command that was executed during the session. |
| <code>vault\sessions\vault_name\sidxxx\detailed_distlist_vault.rpt</code> | The Detailed Distribution List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\distlist_robot.rpt</code> | The Distribution List for Robot report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\distlist_vault.rpt</code> | The Distribution List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\duped.images</code> | List of images successfully duplicated during the session. |
| <code>vault\sessions\vault_name\sidxxx\duplicate.log.nn</code> | Contains output from <code>bpduplicate</code> . |
| <code>vault\sessions\vault_name\sidxxx\eject.list</code> | List of media to be ejected for the session. |
| <code>vault\sessions\vault_name\sidxxx\image.list</code> | NetBackup image catalog information for each image duplicated. |
| <code>vault\sessions\vault_name\sidxxx\image.list_suspend</code> | NetBackup image catalog information for each image whose media will be suspended. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|---|--|
| <code>vault\sessions\vault_name\sidxxx\imagefull.list</code> | A list of all images in the NetBackup image catalog. |
| <code>vault\sessions\vault_name\sidxxx\lostmedia.rpt</code> | The Lost Media Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\media.list</code> | NetBackup media used for originals and duplicates. |
| <code>vault\sessions\vault_name\sidxxx\media.list_suspend</code> | NetBackup media used for originals to be suspended. |
| <code>vault\sessions\vault_name\sidxxx\nb_media.list</code> | Contains the number of images, size of images, and expiration dates for original and duplicated media. |
| <code>vault\sessions\vault_name\sidxxx\nb_media.list_suspend</code> | Contains the number of images, size of images, and expiration dates on original media to be suspended. Used for reports. |
| <code>vault\sessions\vault_name\sidxxx\nbudb.media</code> | Contains media IDs that were allocated for NetBackup database backups. |
| <code>vault\sessions\vault_name\sidxxx\non_vaulted.rpt</code> | The Non-vaulted Images Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\notvaulted.list</code> | A list of images that were not vaulted. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|---|--|
| <code>vault\sessions\vault_name\sidxxx\offsite_inventory.rpt</code> | The Off-site Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\picklist_robot.rpt</code> | The Picking List for Robot report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\picklist_vault.rpt</code> | The Picking List for Vault report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\preview.list</code> | A list of all images that will be considered for duplication or ejection by the current vault session. |
| <code>vault\sessions\vault_name\sidxxx\preview.list_suspend</code> | List of images for which media will be suspended. |
| <code>vault\sessions\vault_name\sidxxx\rcvrimage.inventory</code> | NetBackup image catalog information for all policies between dates specified in the profile. For use with the recovery report. |
| <code>vault\sessions\vault_name\sidxxx\recovery.rpt</code> | The Recovery Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\returned_media.list</code> | Temporary file used for expiring recalled tapes from offsite vault. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|--|---|
| <code>vault\sessions\vault_name\sidxxx\robot.inventory</code> | Lists all the media currently residing in the robot (one media ID per line). |
| <code>vault\sessions\vault_name\sidxxx\summary.log</code> | Concise view of detail.log listing major events during the session, such as how many images were copied. This log is appended for e-mail notification. |
| <code>vault\sessions\vault_name\sidxxx\summary_distlist_vault.rpt</code> | The Summary Distribution List report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\vault.err</code> | Error log for duplication of specific image. |
| <code>vault\sessions\vault_name\sidxxx\vault.err_suspend</code> | Error log for other administrative commands performed during suspend mode; this file should be checked in case of problems. |
| <code>vault\sessions\vault_name\sidxxx\vault.error.file</code> | Error log for other administrative commands performed by Vault; this file should be checked in case of problems. |
| <code>vault\sessions\vault_name\sidxxx\vault_inventory.rpt</code> | The Vault Inventory Report. If a report name includes a time stamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report filenames that include <code>_ff</code> are for debug purposes only; they are created when reports are generated from the Administration Console. |
| <code>vault\sessions\vault_name\sidxxx\volume.db.list</code> | Media Manager inventory for NetBackup database duplicate pool. |
| <code>vault\sessions\vault_name\sidxxx\volume.inventory_suspend</code> | Media Manager inventory for original media pools. |
| <code>vault\sessions\vault_name\sidxxx\volume.list</code> | Media Manager inventory for duplicate pool and NetBackup database duplicate pool. |



Files and Directories for Vault in Windows (continued)

| Directory, Program, or File | Purpose |
|---|---------------------------------------|
| vault\sessions\vault_name\sidxxx\ volume_full.list | Media Manager inventory of all media. |

Vault Functional Design



This functional design document provides an architectural services-oriented approach to building and maintaining a client/server based vault management system. This document defines the functional requirements, develops an architecturally-integrated set of services to meet those requirements, and documents the functional capacity for each service. To provide additional technical detail, this functional design further defines the technical components and technical design needed to provide these services. A final section lists operational procedures needed to deliver each service and assigns basic levels of staff responsibility.

Functional Design Overview

The chief benefit of this functional design is the consistent information for business, technical, and operational viewpoints. High level architectural service design provides a business understanding and real-world value. Technical implementation considerations are shown in the technical component diagrams and definitions. Procedural charts provide a hands-on understanding for operational staff and clarifies areas of responsibility. This section is divided into the following subsections:

- ◆ *Architectural Services:* A set of interrelated services provides a framework for all functional features. The services section uses tables and diagrams to create an organized and expandable system for future enhancements.
- ◆ *Technical Components:* This section explains how the architectural services relate to the actual NetBackup software and/or other software as needed. These diagrams and tables show how services are implemented and interrelate. Tools definitions provide a technical design summarization for architectural considerations, and includes basic interface specifications.
- ◆ *Operational Procedures:* This final subsection shows how the tools are to be used. The procedural diagrams and tables provide a real-world understanding, externalizing any assumptions about who would use the tools and in what manner.



Other Related Services

All storage management services are interdependent in some degree or another. Vault management relies upon both NetBackup and Media Manager services. It is designed to integrate with other operational services subsystems, such as Event Management and Help Desk. The interdependency on these other operational services creates the functional need for close integration of storage management services. Different functional designs provide detailed information for the different storage management sub-systems. It is beyond the scope of this functional design to document these other service areas, such as Help Desk. However, their value to a production environment is noted to provide an operational context.

Other Related Vault Documents

The NetBackup *Vault System Administrator's Guide* provides basic installation, configuration and troubleshooting information. The NetBackup *Vault Operator's Guide* provides day to day procedures to follow to work with vault reports, tapes and vault vendors.

Architectural Services

NetBackup Vault uses a server-based approach that provides both backup duplication and off-site storage and retrieval of media. Vault duplicates backup images onto tape or other media and simplifies restoring the duplicated files when the original backup image media is damaged or unavailable. A master-client implementation extends the storage to other machines by centrally controlling duplication for multiple backup servers simultaneously. A short list of basic services includes:

- ◆ Additional backup protection by duplicating backup images.
- ◆ Optional vaulting of original images/tapes.
- ◆ Direct support for different media types for both backup and duplication media.
- ◆ Backup of images onto different media types to support optimal cost-effective configurations.
- ◆ Maintenance of file operating system level information and security.
- ◆ Scalable, distributed and heterogeneous implementation.
- ◆ Control of tape storage within tape robot and tapes ejected from robot to send to the off-site vault.
- ◆ Assignment of vault slot ID location or container ID when required for use by vault vendor.

- ◆ Appropriate reports
- ◆ Use existing NetBackup capabilities for important functions to ensure compatibility and robustness:
 - ◆ Create catalog backup for Recovery services.
 - ◆ Use Media Manager services for fundamental media and robotic management and control.
 - ◆ Use Media Manager database for keeping track of media containing duplicated images.
 - ◆ Use NetBackup image catalog to keep track of which images need to be duplicated and which images are already duplicated.
 - ◆ Use NetBackup media catalog to determine expiration date of duplicate media.

The overall architectural design is defined as a distributed client/server system.

Client/server systems provide several basic features:

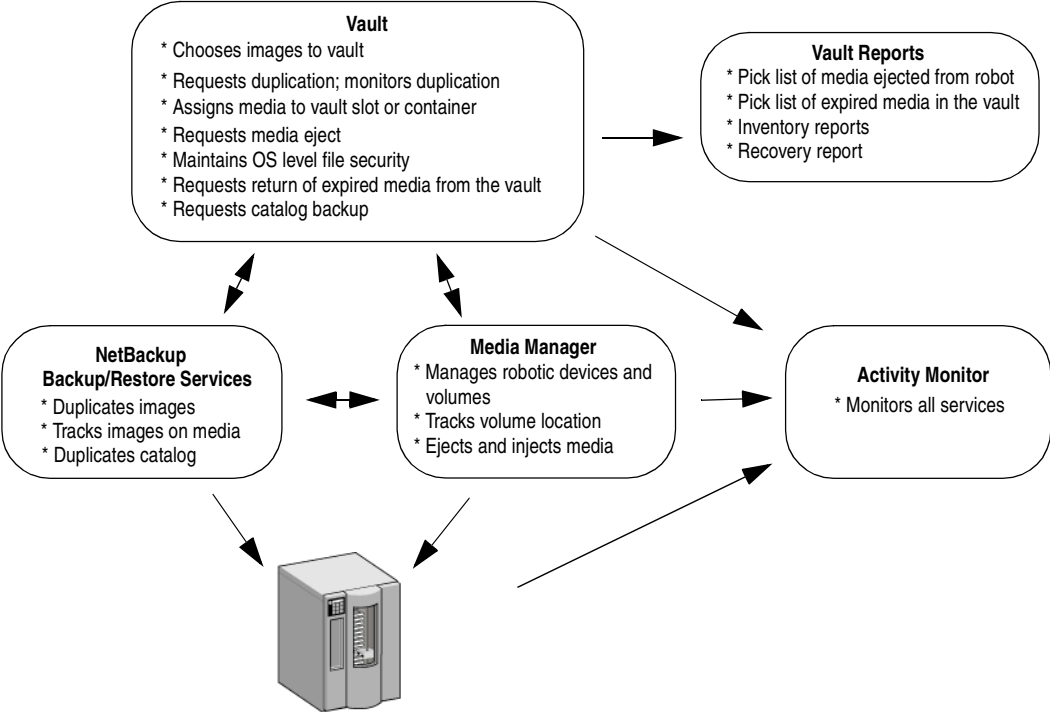
- ◆ Built-in network support. All services must exist within a network topology and thus do not require special configuration or software.
- ◆ Scalable support for large numbers of backup servers. Server based features support faster processors and faster devices.
- ◆ Peer to peer controls. Servers can control other servers to provide better load balancing and different network topologies and bandwidths. Distributed servers ensure better production support and redundancy.
- ◆ Remote installation, control, and configuration. Centralized management reduces management costs by making it easier to setup and run basic backup operations.

Services Interactions Diagram

The following diagram shows the NetBackup, Media Manager, and Vault relationships:



Services Interactions



Client/Server Architectural Services

The following table shows the set of client/server architectural services :

Architectural Services

| Service | Business Challenge | Functional Capability |
|---------------------|--|---|
| Vault Duplication | Protect Information - make sure appropriate images are duplicated Enterprise Scalability - support multiple servers | Determine backup images to duplicate. Duplicate images on multiple drives, multiple servers. Duplicate during day without using production network bandwidth. Optional vaulting of original images without duplication requirement. Can duplicate locally, across LAN or WAN. |
| Vault Monitoring | Fast response | Monitor duplication process for successful completion Support interface to Event Management |
| Vault Configuration | | Set up use of robots, which images to duplicate, other options. |



Architectural Services (continued)

| Service | Business Challenge | Functional Capability |
|-------------------------|---|--|
| Vault Reports | Protect Information - make sure media location is known | Print appropriate reports for sending media off-site to vault and returned from vault. Print regular inventory of vault. |
| Vault Media Management | Reduce Costs - reduce manual administration of media | Keep track of used media location in vault. Eject used media from robot for duplication session. Find returned, expired media and re-use in robot |
| Backup Restore Services | Protect Information | Creates duplicate on image by image basic. De-multiplexes backup image during duplication. Keeps track of both backup images and duplicated copy Keeps track of media used by duplicated copy. Simple image catalog change to restore from copy. |
| Media Manager Interface | Protect Information | Maintains media used information Maintains vault location information. |



Technical Components

The functional requirements outlined in the architectural services section provide you with a general understanding of Vault's capabilities. Additional implementation-specific issues are critical to provide the best quality features. For example, functional scalability creates several technical issues: network bandwidth, catalog sizing, administration, etc. In this section, we list the specific components which deliver the architecture, and review how each component is designed to overcome various technical issues.

Components for Vault

- ◆ Vault Batch Processing
- ◆ Vault Duplication
- ◆ Vault Duplication Monitoring
- ◆ Vault Configuration
- ◆ Vault Reporting
- ◆ Vault Media Management
- ◆ Backup Image Duplication
- ◆ Existing NetBackup Services
- ◆ Media Manager Interface

Technical Design Issues

- ◆ Initiate and control duplications from a centralized location
- ◆ Initiate and control restores of duplicated image from a centralized location
- ◆ Control various duplication parameters by use of NetBackup policies.
- ◆ Reduce duplicating data over the network by duplicating locally whenever possible.
- ◆ Support automated retry of duplication.
- ◆ Support soft shutdown of duplication.
- ◆ Direct appropriate duplications to either master or media servers.
- ◆ Support one or more pairs of drives per server.
- ◆ Write duplicates to a variety of storage devices and media.
- ◆ Allow duplicates to span multi-volume media, yet support industry standard tar format for disaster recovery.

- ◆ Create duplicates that can de-multiplex NetBackup images.
- ◆ Interact with the Media Management service for media availability and media mount/unmount.
- ◆ Create duplicate restores that work the same as normal Backup/Restore, for example, are allowed to the same or a different client, the same or a different location.
- ◆ Allow duplications to notify external operations, for example, by integrating with the Event Management Service.
- ◆ Support all normal backup/restore functions supported, for example: data types, client types.

Vault Technical Components

This table lists the various steps as shown in the Example #1 diagram. Numbers in the table correspond to numbers in the diagram.

Vault Technical Components

| Service | Component | Technical Design |
|---|---|--|
| Vault Batch Functional Capability: Organize various steps into daily/weekly batch | vltrun | Runs Vault utilities for each specific step. Simple script logic Uses one script for ACSLS processing, another for TLD processing. |
| Vault Duplication Functional Capability: ◆ Find all images to duplicate only for policies configured. ◆ Run one or more duplicates simultaneously on one or more servers | Vault - function preview and function duplicate | Run bpimagelist for a given date range. vltrun will filter images based on policies, schedules, schedule types, media servers, and clients. Load balance duplications by splitting found backup images into files, one for each server/drive pair. Run multiple bpduplicates in parallel and monitors bpduplicate processes and log files. Run bpduplicate for each batch of images (based on media ID) found on server where duplicate was made if possible. (If backup server is one of the servers used for duplication.) Otherwise, run duplication on any available server. Can bypass duplication step to only vault original images, original media. |



Vault Technical Components (continued)

| Service | Component | Technical Design |
|--|--|---|
| Vault Duplication Monitoring Functional Capability: Monitor duplicates | vltrun - function duplicate | Monitor log file from bpduplicate for errors in background. Ensure each bpduplicate drive pair output is monitored separately and has unique identifier for support. Ensure bpduplicate has exited before allowing next bpduplicate to run on drive pair. |
| Vault Configuration Functional Capability: Configures which backups to vault | Vault configuration file (db/vault/vault.xml) | The .xml file read by vltrun contains: server(s), number of drive pairs, destination storage unit, policies, schedules, schedule types, clients, media servers, date range, name of vault, NetBackup pools for vaulting and optional duplication, Netbackup robotic and non-robotic group, NetBackup catalog backup pool and retention period. |
| Vault Report Functional Capability: Reports for picking used media from robot/library, for returning expired media from vault, for full vault inventory, for recovery. | vltrun - function report | Reports use data files created by Vault Media Management commands Report "commands" are picking_library, dist_vault, dist_library, picking_vault, vault. |
| Vault Media Management Functional Capability: Commands to find current Media Manager inventory | vltrun - command media | Media function volume inventory - saves Media Manager data for vaulting pools |
| Backup Image Duplication Functional Capability: Find images to duplicate Duplicate images | ◆ bpduplicate | Generates list of images to duplicate Sends command to bptm to copy image, location of log file. Runs in foreground. Waits for bptm to finish. |
| NetBackup Backup Server Functional Capability: ◆ Master server listen for duplication requests ◆ Allow user-directed backups | VERITAS NetBackup Server utilities: ◆ bprd ◆ bpbrm ◆ bptm ◆ bpdm | Support network duplications. Duplicate True Image Recovery (TIR) separately for tracking file deletion. Maintain file / OS security Support multi-tape backup images |



Vault Technical Components (continued)

| Service | Component | Technical Design |
|--|---|--|
| <p>NetBackup Master / Media Server Functional Capability:</p> <ul style="list-style-type: none"> ♦ Master server sends duplicate request to appropriate server ♦ Provide central backup/duplicate image catalog for restore. | <p>VERITAS NetBackup Master/Media Server:</p> <ul style="list-style-type: none"> ♦ <code>bprd</code> ♦ <code>bpcd</code> ♦ <code>add_slave</code> ♦ <code>add_slave_on_clients</code> | <p>Redirect duplication requests to appropriate Media Server based upon destination storage requested for specific image.</p> <p>Allows duplicate to run off-network for tape-to-tape copy.</p> |
| <p>NetBackup Server side Multiplexing (MPX)</p> <p>Functional Capability:</p> <p>Support copy of tape multiplexing both on local backups and across networks.</p> | <p>VERITAS NetBackup server utilities:</p> <ul style="list-style-type: none"> ♦ <code>bpbrm</code> ♦ <code>bptm</code> | <p>De-multiplexing of backup image automatic by <code>bptm</code>, and optionally preserve multiplexing information.</p> <p>Only one duplicate per drive pair.</p> <p>Duplication speed limited to tape read/write speeds. High use of multiplexing limits copy speed to tape read and tape mount for all tapes needed for de-multiplexing, assuming tape read/write speeds similar.</p> |
| <p>NetBackup Backup Server Database Services</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> ♦ Provides Image information including number of copies. ♦ Stores media used information for both primary copy and duplicate copies. | <p>VERITAS NetBackup database daemon/service:</p> <ul style="list-style-type: none"> ♦ <code>bpdbm</code> | <p>Use <code>bpdbm</code> for all data access requests</p> <p><code>bpdbm</code> daemon/service process always running on master server - not on media servers.</p> <p><code>bpdbm</code> creates unique backup identifier based on client name, time of day.</p> <p><code>bpdbm</code> provides image catalog, number of image copies for determining which images need duplication.</p> <p><code>bpdbm</code> provides backup media catalog for determining which backup media contains backup image, and which duplication media is used for copies.</p> <p><code>bpdbm</code> provides restore with duplicate copy image/media when "primary copy" is not 1.</p> |
| <p>NetBackup Schedules - Centralized server scheduling</p> <p>Functional Capability:</p> <p>Can be used by <code>bpduplicate</code> to limit duplication of a policy to only images within specific schedule</p> | <p>VERITAS NetBackup Scheduler:</p> <ul style="list-style-type: none"> ♦ <code>bpsched</code> | <p><code>vltrun</code> preview can be limited by schedule name or type.</p> |
| <p>NetBackup Backup Server - Restore functions</p> <p>Functional Capability:</p> <p>Works as normal for duplicated images.</p> | <p>VERITAS NetBackup server daemons/processes:</p> <ul style="list-style-type: none"> ♦ <code>bprd</code> ♦ <code>bpcd</code> ♦ <code>bpbrm</code> ♦ <code>bptm</code> ♦ <code>bpdm</code> | <p>Duplicated images can be de-multiplexed.</p> <p>Changing the primary copy automatically forces any restore of that image to use the specified primary copy.</p> <p>There is currently no graphic user interface to select the duplicated copy. The change requires manual intervention.</p> |



Vault Technical Components (continued)

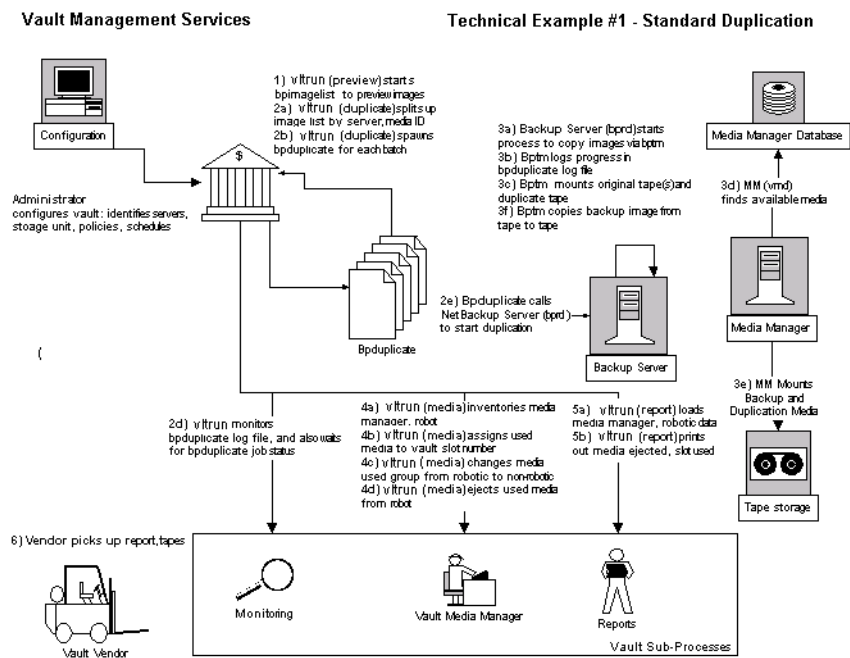
| Service | Component | Technical Design |
|--|--|--|
| <p>NetBackup policies - Centralized client management</p> <p>Functional Capability:</p> <p>Works as normal. Policy name used to limit duplication to specific clients.</p> | <p>VERITAS NetBackup policy utilities</p> | <p>Name of clients stored in the configuration file.</p> |
| <p>NetBackup Storage Unit - Centralized resource control</p> <p>Functional Capability:</p> <p>Works as normal. Used to determine which server will run duplicate</p> | <p>VERITAS NetBackup storage unit utilities</p> | <p>Destination storage unit in configuration file will match up with source media server.</p> |
| <p>NetBackup Image Catalog Services - centralized image information</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> Centrally stores backup and duplicate catalog information. Keeps track of both backup and duplicate image information: media, fragments, size, location on tape required for recovery. Lists directories/files/fragments in image catalog. Tracks backup and duplicate server name for each image to support master/media server duplication | <p>VERITAS NetBackup image catalog utilities:</p> <ul style="list-style-type: none"> bptm bpimage bpimagelist bplist bpflist bpimmedia bpfrag | <p>The master server stores all image data. Duplicates notify the master server of copy status - fragments, media used, but no image information necessary.</p> |
| <p>NetBackup Media Manager - centralized media control</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> Used to determine which media used for primary backup copy. Can manually expire primary copies of images to force use of duplicate copy. Backup retention periods for primary backup image can differ for duplicate image. Expire media after retention period will expire duplicate copies. Allow expired frozen media sent off-site to remain off-site while needed for recovery. | <p>VERITAS NetBackup media database utilities:</p> <ul style="list-style-type: none"> bptm bpmedia bpmedialist bpexpdate bplabel | <p>Media used for image duplicates can have a different retention period.</p> <p>Duplicate media used must be assigned in Media Manager to duplicate pool named in configuration file.</p> <p>Expired duplicate media released in Media Manager triggers recall from off-site vault.</p> <p>Duplication sends image fragment, media used information to master server.</p> |

Vault Technical Components (continued)

| Service | Component | Technical Design |
|--|--|--|
| Backup monitoring Functional Capability: <ul style="list-style-type: none">◆ Duplication monitored by NetBackup monitoring tools.◆ Vault log files can be used to track status. | Activity Monitor - VERITAS NetBackup daemon/service logging. | VERBOSE option increases information to logging subdirectories. Notify scripts supported for start and end session, and for pre-and post- eject processing. |
| Backup reports Functional Capability: Provide administrative reports needed for capacity, utilization planning, auditing purposes, technical support. | See Backup Console. <ul style="list-style-type: none">◆ bpererror◆ bpimagelist◆ cleanstats◆ available_media◆ support | Can load backup information into RDBMS for better reporting. bpimagelist provides basic information showing all copies of an image. |
| Media Manager interface Functional Capability: <ul style="list-style-type: none">◆ Use Media Manager Tools to request media from specific pools; mount tapes; control robotics◆ Use specific Media Manager pool for duplication. | | Must ensure NetBackup Media database and Media Manager database stay synchronized. Update Media Manager fields to store vault name, duplication session, slot id, date requested, date sent off site. |



Technical Example: Standard Duplication Diagram



Technical Example: Standard Duplication Table

Technical Example #1: Standard Duplication in Vault

| Service | Component | Outgoing Program/Data Flow |
|--|------------------------------------|--|
| Backup Server-File System and Raw Partition Incoming Program/Data Flow: | bprd - Job request daemon/service. | bprd starts bpsched on master server via command line interface to handle jobs. |
| ♦ bprd started by startup script on master server via command line. bprd must always be running to allow any backup/restore command. | | For manual jobs, bprd builds file list from client (for example, bpbbackup), then bprd starts bpsched as normal. |
| ♦ bprd can be started in the Administration Console Activity Monitor | | bprd regularly cleans up debug logs. |
| ♦ bprd called by bpbbackup, bpduplicate, bprestore, bparchive to start jobs. Uses known socket in /etc/services | | |
| ♦ Debug Log: /usr/openv/netbackup/logs/bprd on server | | |



Technical Example #1: Standard Duplication in Vault (continued)

| Service | Component | Outgoing Program/Data Flow |
|--|---|---|
| <p>Schedules - Centralized backup scheduling</p> <p>Incoming Program/Data Flow</p> <p>Scheduler started by <code>bprd</code> on master server.</p> <p>Scheduler exits if no jobs needed to run, or monitored client job is finished.</p> <p>Debug Log: <code>/usr/openv/netbackup/logs/bpsched</code> on Master server</p> | <p>Master Server Scheduler:</p> <p><code>bpsched</code></p> | <p><code>bpsched</code> calls <code>bpdbm</code> to obtain policy information, storage unit server to use, file list, etc.</p> <p><code>bpsched</code> starts <code>bptm</code> to do basic media check prior to starting client job.</p> <p><code>bpsched</code> starts <code>bpcd</code> via <code>inet .d</code> on appropriate storage unit server (master or media server).</p> <p><code>bpsched</code> sends <code>bpbrm</code> command for <code>bpcd</code> to run, with required options, file list, log path to <code>bpcd</code> on server via socket</p> <p><code>bpsched</code> monitors <code>bpbrm</code> output via <code>stderr</code></p> <p><code>bpsched</code> signals other <code>bpsched</code> processes to ensure only one scheduler is acting as main scheduler, to eliminate accidental duplication of jobs.</p> |
| <p>See Example #1 Diagrams</p> <p>Incoming Program/Data Flow</p> <p><code>bpcd</code> started by <code>bpsched</code> via <code>inetd</code> / known socket.</p> <p>Acts as job proxy for <code>bpsched</code></p> <p>Debug Log: <code>/usr/openv/netbackup/logs/bpcd</code> on server for server related commands</p> | <p><code>bpcd</code></p> <p>Backup job proxy</p> | <p><code>bpcd</code> starts backup/restore manager <code>bpbrm</code> on master or media server</p> |
| <p>Backup Server (continued)</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> ◆ <code>bpbrm</code> started by <code>bpsched</code> on storage unit server (via <code>bpcd</code>). One <code>bpbrm</code> process started for each backup or restore operation. ◆ <code>bpbrm</code> <code>stderr</code> output and exit status monitored by <code>bpsched</code> ◆ <code>bpbrm</code> exits to <code>bpsched</code> with status ◆ <code>bpbrm</code> receives <code>signal()</code> from <code>bptm</code> when media ready. ◆ <code>bpbrm</code> sends e-mail on job completion (via <code>bpcd</code>) <p>Debug Log: <code>/usr/openv/netbackup/logs/bpbrm</code> on server</p> | <p><code>bpbrm</code></p> <p>Backup/Restore Manager</p> | <p><code>bpbrm</code> starts <code>bptm</code> to start backups and duplication. Waits for child exit status.</p> <p><code>bpbrm</code> starts <code>bpbkarr</code> on client. Waits for child exit status.</p> |



Technical Example #1: Standard Duplication in Vault (continued)

| Service | Component | Outgoing Program/Data Flow |
|---|---|--|
| <p>Backup Server (continued)</p> <p>Incoming Program/ Data Flow:</p> <ul style="list-style-type: none"> ♦ bptm started by bpsched to check validity of storage unit. ♦ bptm started by bpbrm on master and Media Servers via command line. One bptm started for backup/restore when using tape or optical media. ♦ Child bptm started by Parent bptm. ♦ bptm messages bpbrm once media is mounted via signal. ♦ bptm stderr output monitored by bpbrm ♦ Child bptm exits to parent bptm on client completion. ♦ Parent bptm exits to bpbrm on server completion. <p>Debug Log (on server): UNIX: /usr/opensv/netbackup/logs/bpsched</p> <p>Windows: install_path/netbackup/logs/bpsched</p> | <p>bptm: Server tape manager</p> | <p>Parent bptm starts child bptm.</p> <p>Parent bptm calls vmd on master server via known socket to find appropriate media.</p> <p>Parent bptm also calls bpdgm to compare Media Manager database with own, NetBackup Media database.</p> <p>Parent bptm calls ltid on storage unit server via known socket to mount media.</p> <p>Child bptm receives data from bpbkar on client.</p> <p>Child bptm writes data to buffer.</p> <p>Parent bptm writes buffers to tape when they are full.</p> <p>On backup completion, parent bptm runs:</p> <p>UNIX: /usr/opensv/netbackup/bin/backup_notify</p> <p>Windows: install_path/netbackup/bin/backup_notify</p> |
| <p>Backup Server Policy, Image and Media Database Services</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> ♦ bpdgm started by initbprd during OS boot on master server only. ♦ In this example, bpdgm called by bpsched, bpbrm, bptm. ♦ bpdgm daemon/service process always running on master server - not on media servers. ♦ bpdgm database can be accessed locally or across network using known socket for bpdgm. ♦ bpdgm creates unique backup identifier based on client name, time of day. | <p>Master Server database daemon/service: bpdgm</p> | <p>bpdgm provides policy data to bpsched to build worklist.</p> <p>bpdgm provides configuration information to bpbrm.</p> <p>bpdgm provides bptm with information about Backup Media database for comparison with Media Manager database.</p> <p>bpdgm stores Media used information into Backup Media database for bptm</p> <p>bpdgm stores Image file list catalog for bpbrm.</p> |
| <p>Backup Media Manager</p> <p>Incoming Program/Data Flow:</p> <p>Backup Media management handled by calls from bptm to bpdgm shown above.</p> | <p>bptm bpdgm</p> | <p>Parent bptm calls bpdgm to save media used data to media database. Updates keyword search for user-directed backups.</p> |
| <p>Backup Image catalog Incoming Program/Data Flow:</p> <p>Image catalog handled by calls from the client bpbkar to server bpbrm and then from bpbrm to bpdgm.</p> | <p>bpbkar bpbrm bpdgm</p> | <p>bpbkar writes image catalog data to bpbrm on server.</p> <p>bpbrm calls bpdgm via socket to store image file list catalog.</p> |



Technical Example #1: Standard Duplication in Vault (continued)

| Service | Component | Outgoing Program/Data Flow |
|--|--------------------------------------|--|
| <p>Backup Client</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> ◆ bpcd started by bpbrm (via inetd) to start backup jobs. Acts as proxy for bpbrm. ◆ bpcd can perform other chores, e.g. sends mail for bpbrm. ◆ bpcd returns exit status to bpbrm via socket. <p>Debug Log:</p> <p>UNIX: /usr/opensv/netbackup/logs/bpcd</p> <p>Windows: <i>install_path</i>/netbackup/logs/bpcd</p> | <p>bpcd - client job proxy</p> | <p>bpcd starts bpbkar on client for backup.</p> |
| <p>Backup Client, continued</p> <p>Incoming Program/Data Flow:</p> <p>bpbkar started by bpbrm (via bpcd).</p> <ul style="list-style-type: none"> ◆ bpbkar receives backup policy options, file list, etc. from bpbrm. This information was originated by bpsched at step #1. ◆ bpbkar exit status is sent to bpsched via socket. <p>Debug Log:</p> <p>UNIX: /usr/opensv/netbackup/logs/bpbkar</p> <p>Windows: <i>install_path</i>/netbackup/logs/bpbkar</p> | <p>bpbkar - client data transfer</p> | <p>bpbkar reads file/directory list from bpbrm.</p> <p>bpbkar determines NFS mounts and adds or ignores depending on policy.</p> <p>bpbkar compresses client data if necessary.</p> <p>bpbkar writes client data to bptm/bpdm on server.</p> |
| <p>Media Manager</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> ◆ Media Manager vmd called by server tape manager bptm for scratch tape. ◆ Media manager ltid called by server tape manager bptm to mount tape. <p>Note See Media Manager functional design for detailed information.</p> | <p>vmd</p> <p>ltid</p> | <p>Media Manager provides new media ID for backup.</p> <p>Uses pool for determining which media is valid.</p> <p>Media Manager ltid mounts tape.</p> |



Operational Procedures

This table summarizes operational procedures for Vault. The *NetBackup Vault Operator's Guide* provides more detailed information on day-to-day procedures. The *NetBackup Vault System Administrator's Guide* provides more detailed information on installation, configuration and troubleshooting.

Vault Management / Operational Procedures

| Service | Operational Procedure | Staff Responsibilities |
|----------------------------|---|--|
| Vault Configuration | Review backup procedures; determine duplication capacity needed. Assign appropriate server to run duplications; determine appropriate window for running duplication. Configure Vault parameters via Administration Console. Review duplication windows for performance, throughput. | Determine need for basic levels of duplication service on a per policy basis. Ensure sufficient hardware, software, network capacity is available for duplication of backup images. |
| Vault Duplication | Set up Vault policy to schedule vault sessions. | Start duplication job on time, daily and/or weekend. |
| Vault Monitoring | Use Activity Monitor to determine progress. Set up links between log file and monitoring system for e-mail and/or paging notification. | Ensure duplication jobs complete successfully. Ensure errors reported to Event Management. |
| Vault Report | Compare report output with ejected tapes, returned tapes | Ensure accuracy of vault process. |
| Vault Media Management | Check duplication volume pools and catalog backup pools for sufficient media. | Ensure sufficient media available for duplication. |
| Backup Media Manager | Use Media Manager to manage media. Manually expire/freeze tapes when needed for retrieval from vault. | |
| Backup Image Catalog | Set up schedule for backup of image catalog. Ensure specific tapes available to store catalog. | Ensure backup catalog is safely copied. |
| Event Management Interface | Set up appropriate Event response procedures. | Same as for normal Backup. |



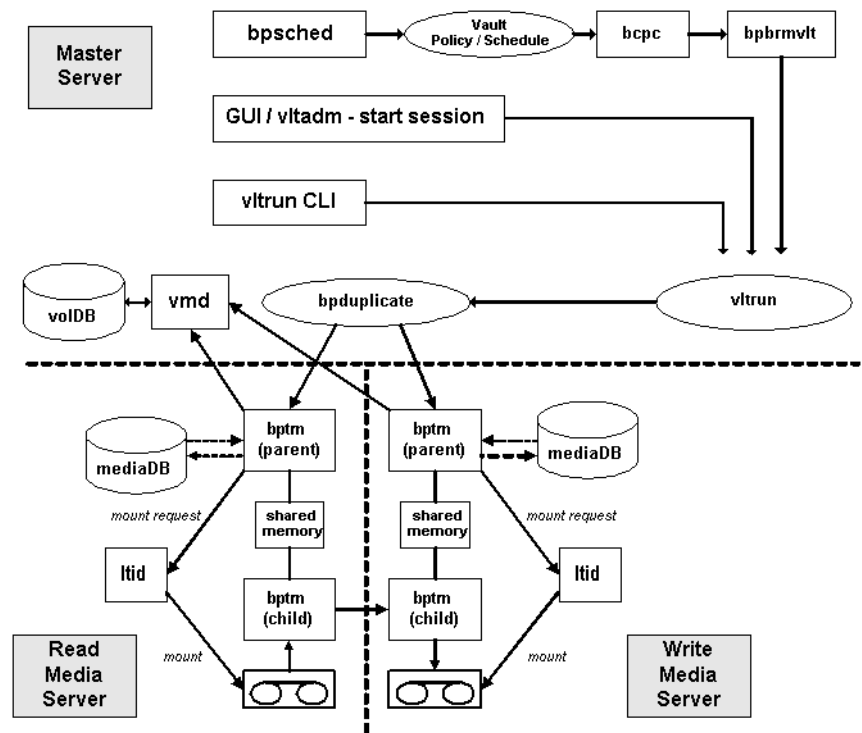
Vault Management / Operational Procedures (continued)

| Service | Operational Procedure | Staff Responsibilities |
|-------------------------------|--|--|
| Incident Management Interface | Set up procedure for passing storage events to Help Desk. | Same as for normal Backup. |
| Backup Reports | Run regular reports to ensure proper images are duplicated. | Review production duplication cycle for thoroughness. |
| Duplication Capacity Review | Determine capacity planning cycle, including reaction time, costing factors, and new requirements. | Assist production support over time on determining system, robotic, network utilization rates and effective valuation, for example, disk capacity. Assist in new requirements and performance-related additions to the system infrastructure. |
| Recovery Review | Run regular tests to ensure recovery of essential data from off-site storage. | Ensure knowledge of appropriate procedures for restoring duplicated images. Ensure sufficient knowledge to restore database catalog, backup software, etc. in case of disaster on Netbackup server(s). |
| Media Manager | Determine media requirements and setup initial media pool for duplication. Monitor media pool usage over time. Configure master/media server media management. | Ensure sufficient media is available for duplicates to run. |

NetBackup Vault Image Duplication Process

The diagram below shows the flow of operations for the duplication process during a vault session. For simplicity, this diagram shows duplication from tape-to-tape, and shows operations for creating a single copy.





A vault session is initiated by:

- ◆ The NetBackup scheduler (bpsched) finding a Vault schedule with an open window.
- ◆ The operator issuing a Start Session for a particular Vault profile from the NetBackup Administration Console >Vault Management or from the Vault Administration (vltadm) menu.
- ◆ The operator issuing a vltrun command from the command line.

Regardless of the method that initiated the vault session, it will result in vltrun being started. These are the stages of the process:

Image Duplication Process

| Process | Actions |
|---------|--|
| vltrun | vltrun starts bpduplicate. vltrun performs pre-processing to determine which images need to be duplicated, ejected, and so on. |

Image Duplication Process (continued)

| Process | Actions |
|----------------------|--|
| bpduplicate | bpduplicate starts bptm on the READ SERVER and on the WRITE SERVER. |
| bptm | bptm on the READ SERVER and on the WRITE SERVER respectively locate media to be used in the duplication process (lookups to mediaDB and volDB) and issue mount requests (tpreq) for those media. |
| bptm (child process) | bptm (child process) read the data from the READ SERVER media, send it to the WRITE SERVER where the bptm (child process) writes the data to the destination media. |



Index

A

- access management 180
- accessibility features xx
- ACS number 51
- ACSLs server 51
- adding alternate media server names 23, 53
- administering access to Vault 180
- advanced duplication 68
 - to avoid sending data over the network 36
- All Media Inventory report 172
- alternate media server names
 - adding 23, 53
- alternate read server
 - to avoid sending data over the network 35
- Alternate Read Server setting 70, 77
- Any Available storage unit setting 31, 32, 33
- assigning multiple retentions with one profile 121
- At Time of Eject setting (Suspend This Session's Media) 87

B

- Backup Policies setting 64
- Backup Policy For Multiple-Tape Catalog
 - Backups setting 84
- Backups Started setting 65
- best practices 19
- bpdjobs
 - changes for Vault 195

C

- CAP. *See* media access port
- catalog backup
 - Backup Policy for Multiple-tape Catalog
 - Backups setting 84
 - ensuring media for 134
 - Files To Be Backed Up window 84

- Media Server setting 84
- Retention Period setting 85
- reusing unexpired catalog backup media 135
- Skip the Catalog Backup Step setting 85
- Use Catalog File Locations from NetBackup Catalog Backup Configuration setting 85
- Volume Pool setting 85
- Catalog Backup tab 82
- changing off-site volume group name 184
- changing off-site volume pool name 184
- changing volume group name 184
- changing volume pool name 184
- choose backups
 - Backup Policies setting 64
 - Backups Started setting 65
 - Clients setting 65
 - Media Servers setting 65
 - Schedules setting 65
 - Source Volume Group setting 66
 - Type Of Backups setting 66
- Choose Backups tab 63
- clearing the media description field 134
- clearing the vault fields 134
- Clients setting 65
- clusters 7
 - shared disk 218, 225
- Complete Inventory List for Vault. *See* All Media Inventory report
- concurrent copies
 - continue or fail 144
 - during advanced duplication 155
 - during basic duplication 151
 - during original backup 145
 - during Vault duplication 147
 - overview 143
 - through the NetBackup Catalog



- node 149
- to avoid sending duplicates over the network 35
- configuration
 - Catalog Backup tab 82
 - Choose Backups tab 63
 - Duplication tab 66
 - Eject tab 85
 - methods 51
 - profile 62
 - Reports tab 91
 - robots 56
 - volume pools 43
- Container Inventory Report 173
- Containers of Many Media setting 59
- continue
 - for concurrent copies 144
- Copies setting 75, 78
- copying a profile 182
- creating
 - a profile 61
 - a vault 57
 - a vault policy 47
- Customer ID setting 59

D

- debug logs 202
- deferred reports 93
- Deferred setting (Eject Mode) 87
- Destination Storage Unit setting 73, 76, 80
- Destination Volume Pool setting 74, 76, 80
- Detailed Distribution List for Vault 166
- Directory structure 217
- directory structure 217
- disaster recovery
 - definition 206
 - definition of disaster recovery plan 207
 - developing a disaster recovery plan 207
 - preparing for disaster 205
 - priorities 207
 - testing a disaster recovery plan 209
- disk staging 21
- Distribution List for Vault 165, 169
- duplicate images 2
- duplication
 - advanced 68
 - Alternate Read Server setting 70, 77
 - basic 67
 - Copies setting 75, 78

- Destination Storage Unit setting 73, 76, 80
- Destination Volume Pool setting 74, 76, 80
- Expire Original Disk Backup Images setting 71
- Fail Copies setting 75, 78
- increase throughput 37
- Multiple Copies dialog 74
- Multiple Copies setting 71
- multiplexed 148
- Number Of Read Drives setting 72, 78
- Number Of Write Drives setting 74, 76, 80
- Preserve Multiplexing setting 72
- Primary Copy setting 71, 75, 79
- Retention Level setting 73
- Retention setting 76, 79
- Source Backup Server setting 78
- Source Backups Reside On setting 73, 79
- through the NetBackup Catalog
 - node 149
 - when possible 148
- Duplication tab 66, 68
- duplication throughput
 - configuring for multiple drives 38
 - multiple-drive scenario 38

E

- eject
 - consolidating ejects 111
 - Eject Mode setting 87
 - Suspend Media On Which Backups Were Written setting 88
- Eject Mode settings 87
- Eject tab 85
- ejected tapes returned to robot 201
- ejecting partial images
 - use suspend to avoid 24
- e-mail
 - notifying a tape operator when eject begins 130
- error codes 198
 - extended 103
- EXIT status 198
- Expire Original Disk Backup Images setting 71
- extended error codes 103



F

- fail
 - for concurrent copies 144
- Fail Copies setting 75, 78
- files and directories 217
- Files To Be Backed Up window 84
- First Off-Site Slot ID setting 60
- Full Inventory List for Vault. *See* Off-site Inventory report

G

- Glossary. *See* NetBackup Help.

I

- images
 - duplicate 2
 - original 2
 - primary backup 67
- Immediate Eject setting 87
- immediate reports 94
- Immediate setting (Eject Mode) 87
- Immediately setting (Suspend This Session's Media) 87
- Inline Tape Copy. *See* concurrent copies
- installation
 - on UNIX systems 8
 - on Windows systems 15
 - prerequisites for a UNIX system 8
 - prerequisites for a Windows systems 16
- Inventory List for Vault. *See* Vault Inventory report
- Iron Mountain Electronic Format report 176

L

- license key
 - adding on Windows systems 16
- load balancing
 - by duplicating daily and ejecting weekly 34
 - profiles for both originals and duplicates 34
- log files
 - debug logs 202
 - set duration 203
 - vault session 184
- LSM number 51

M

- MAP. *See* media access port
- master server
 - host name of 50

media

- for catalog backup 134
- media access port
 - capacity 51
 - number 51
 - to use for eject 60
- media description field
 - clearing 134
- media missing in robot 198
- media server names 50
 - adding alternate 23, 53
- Media Server setting for catalog backup 84
- Media Servers setting 65
- menu user interface
 - bpdjobs 195
 - changes in vmadm 192
 - overview 189
 - Vault Administration Interface 189
 - Vault Oerator Menu Interface 191
 - vltadm 189
 - vltopmenu 191
- Monitoring a vault session 102
- moving a vault to different robot 183
- multiple copies
 - overview 143
 - See Also* concurrent copies
- Multiple Copies dialog 74
- Multiple Copies setting 71
- multiple retention mappings 121
- multiple volume groups 23
- multiplexed duplication 148

N

- network
 - avoid sending duplicates over the network 35
- New Profile dialog 62
- New Vault dialog 59
- New Vault Robot dialog 56
- notify scripts
 - for a specific profile 133
 - for a specific robot 132
 - for a specific vault 133
 - order of execution 133
 - using 131
 - vlt_ejectlist_notify 132
 - vlt_end_notify 132
 - vlt_endeject_notify 132
 - vlt_start_notify 131



- vlt_starteject_notify 132
- notifying a tape operator when eject begins 130
- Number Of Read Drives setting 72, 78
- Number Of Write Drives setting 74, 76, 80

O

- Off-site Inventory report 171
- off-site volume group 45, 60
 - changing name of
 - off-site volume group
 - renaming 184
- Off-Site Volume Group setting 60
- off-site volume pool 43
 - changing name of 184
 - renaming 184
- Off-Site Volume Pools setting 87
- Off-site Volume Pools windows (Eject tab) 87
- organizing reports by profile 41
- organizing reports by robot 41
- organizing reports by vault 41
- original images 2

P

- partial backups
 - use suspend to avoid vaulting 24
- Picking List for Robot 164
- Picking List for Vault 168
- policy 47
 - configuration information 46
 - names 46
 - NetBackup policy for Vault 46
 - schedule names 46
- preferred vaulting strategies 20
- Preserve Multiplexing setting 72
- preview media to be ejected 106
- preview vault session 100
- primary backup image 67
- Primary Copy setting 71, 75, 79
- printing
 - troubleshooting problems 197
- profile
 - configuring 62
 - copying a 182
 - creating 61
 - notify script for a specific 133
 - organizing reports by 41
 - overlap time window 22
 - printing information 182

- profiles
 - for both originals and duplicates 34

R

- recovering backup images 137
- recovery
 - keep primary copy on site 27
 - match volume pools to usage 27
 - of damaged media 137
 - preparing for efficient 26
 - revault unexpired media 29
 - use precise naming conventions 27
 - vault NetBackup catalogs 26
- Recovery Report for Vault 174
- renaming off-site volume group 184
- renaming off-site volume pool 184
- renaming volume group 184
- renaming volume pool 184
- reports
 - All Media Inventory 172
 - Complete Inventory List for Vault. *See* All Media Inventory report
 - consolidating reports 162
 - Container Inventory Report 173
 - deferred 93, 94
 - Detailed Distribution List for Vault 166
 - Distribution List for Robot 169
 - Distribution List for Vault 165
 - Full Inventory List for Vault. *See* Off-site Inventory report
 - immediate 94
 - inventory 170
 - Inventory List for Vault. *See* Vault Inventory report
 - Iron Mountain Electronic Format 176
 - Lost Media Report 41, 175
 - Off-site Inventory 171
 - organizing by profile 41
 - organizing by robot 41
 - organizing by vault 41
 - Picking List for Robot 164
 - Picking List for Vault 168
 - printing 159
 - Recovery Report for Vault 174
 - Summary Distribution List for Vault 167
 - types of 164
 - Vault Inventory report 170
- Reports tab 91
- resource contention



- avoid by robot usage 30
- avoiding 29
- load balancing 34
- sharing resources with backup jobs 33
- specifying different volume pools for source and destination 35
- use disk staging to avoid 21
- vault original backups to avoid 21
- when the read drive is not in the vault robot 33
- when two processes try to use the same drive 30
- resuming a vault session 101
- Retention Level setting 73, 76, 79
- retention period based on copy one
 - retention period 121
- Retention Period setting for catalog backups 85
- retention_mappings file 123
- revault media 126
- robot
 - configuring for vault 56
 - control host 57
 - notify script for a specific 132
 - number 57
 - organizing reports by 41
 - properties 51, 57
 - types of 50
- Robot Name setting 57
- Robot Number setting 57
- Robot Type setting 57
- robotic volume group 23, 45
- Robotic Volume Group setting 60
- running multiple sessions
 - simultaneously 98

S

- Schedules setting 65
- scratch volume pools 40
- server name group 53
- session files
 - setting the duration of 186
- setting the duration of session files 186
- Skip the Catalog Backup Step setting 85
- Skip the Eject Step setting 87
- Slot ID 60
- Slots for Individual Media setting 60
- Source Backup Server setting 78
- Source Backups Reside On setting 73, 79

- source volume group 24, 32
- Source Volume Group setting 66
- status codes 198
- storage unit
 - Any Available setting 31, 32, 33
 - name of 50
 - number of drives in 50
- Summary Distribution List for Vault 167
- supported clusters 7
- supported robots 7
- supported systems 7
- suspend
 - Suspend Media On Which Backups Were Written setting 24
 - Suspend Option setting 84
 - use suspend to avoid partial backups 24
- Suspend Media for the Next Session setting 88
- Suspend Media on Which Backups Were Written setting 88
- Suspend Option setting 88
- Suspend This Session's Media setting 88

T

- time window
 - overlapping 22
- troubleshooting
 - bad or missing duplicate media 199
 - drive or robot offline 200
 - ejected tapes returned to robot 201
 - ejecting tapes while in use 201
 - error codes 198
 - logs 204
 - media missing in robot 198
 - no duplicate progress message 200
 - printing problems 197
 - session locking 202
- Types Of Backups setting 66

U

- unassigning catalog backup media 135
- uninstall Vault
 - from UNIX systems 13
- UNIX files and directories 217
- Use NetBackup Catalog Paths setting 85

V

- Vault
 - accessing 2
- vault



-
- creating a 57
 - moving to a different robot 183
 - name 61
 - notify script for a specific 133
 - only the intended backups 24
 - organizing reports by 41
 - original backups 21
 - paradigm 2, 20
 - policy 46
 - printing vault information 182
 - vendor 61
 - Vault Administration Interface 189
 - Vault Inventory report 170
 - Vault Name setting 61
 - Vault Operator Menu Interface 191
 - vault original backups 21
 - vault session
 - locking 202
 - previewing 100
 - resuming 101
 - running 98
 - running multiple sessions
 - simultaneously 98
 - setting the duration of session files 186
 - stopping 101
 - Vault Vendor setting 61
 - vaulting
 - containers 59
 - defined 2
 - original backups in a 24x7
 - environment 25
 - paradigm 2, 20
 - preferred strategies 20
 - vmadm
 - changes for Vault 192
 - special actions menu 192
 - volume configuration 192
 - Volume Database Host setting 57
 - volume group
 - changing name of 184
 - configuration 45
 - off-site 45
 - overview 4
 - renaming 184
 - robotic 45, 60
 - source 66
 - volume pool
 - changing name of 184
 - configuration 43
 - destination 74, 76, 80
 - for backup media that remains on site 43
 - for catalog media 44
 - match to data usage 27
 - naming conventions 27
 - off-site 43
 - overview 5
 - renaming 184
 - use scratch volume pools 40
 - Volume Pool setting for catalog backup 85
 - Volume pools
 - specifying different source and destination 35
- W**
- Windows files and directories 224
 - Windows systems
 - delicensing Vault 17
 - working files
 - setting the duration of 186

